

IAP/ Rec'd PCT/PTO 11 JUL 2006

PTO/SB/21 (09-04)

Approved for use through 07/31/2006. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL
FORM

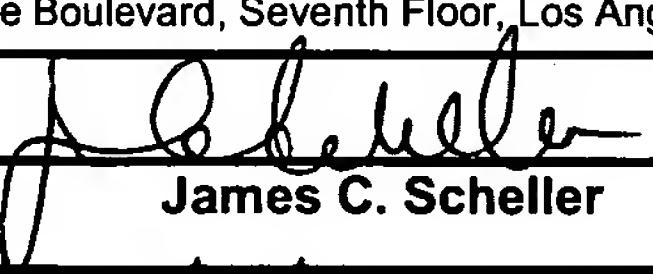
(to be used for all correspondence after initial filing)

		Application Number	10/580,557
		Filing Date	26 May 2006
		First Named Inventor	Benjamin Liu
		Art Unit	Not Yet Assigned
		Examiner Name	Not Yet Assigned
Total Number of Pages in This Submission	39	Attorney Docket Number	42390P22191

ENCLOSURES (Check all that apply)

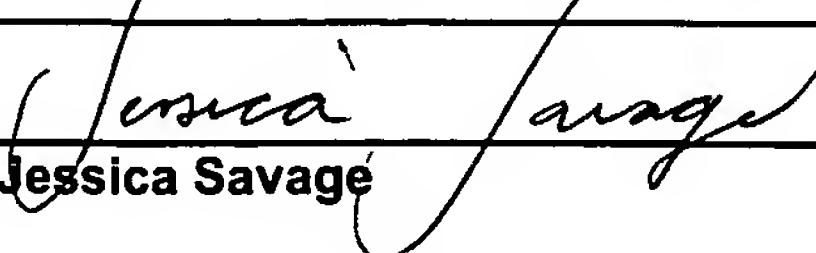
<input type="checkbox"/> Fee Transmittal Form	<input type="checkbox"/> Drawing(s)	<input type="checkbox"/> After Allowance Communication to TC
<input type="checkbox"/> Fee Attached	<input type="checkbox"/> Licensing-related Papers	<input type="checkbox"/> Appeal Communication to Board of Appeals and Interferences
<input type="checkbox"/> Amendment/Reply	<input type="checkbox"/> Petition	<input type="checkbox"/> Appeal Communication to TC (Appeal Notice, Brief, Reply Brief)
<input type="checkbox"/> After Final	<input type="checkbox"/> Petition to Convert to a Provisional Application	<input type="checkbox"/> Proprietary Information
<input type="checkbox"/> Affidavits/declaration(s)	<input type="checkbox"/> Power of Attorney, Revocation	<input type="checkbox"/> Status Letter
<input type="checkbox"/> Extension of Time Request	<input type="checkbox"/> Change of Correspondence Address	<input checked="" type="checkbox"/> Other Enclosure(s) (please Identify below): Return Receipt Postcard
<input type="checkbox"/> Express Abandonment Request	<input type="checkbox"/> Terminal Disclaimer	
<input type="checkbox"/> Information Disclosure Statement	<input type="checkbox"/> Request for Refund	
<input checked="" type="checkbox"/> Certified Copy of Priority Document(s)	<input type="checkbox"/> CD, Number of CD(s) _____	
<input type="checkbox"/> Reply to Missing Parts/ Incomplete Application	<input type="checkbox"/> Landscape Table on CD	
<input type="checkbox"/> Reply to Missing Parts under 37 CFR 1.52 or 1.53		
Remarks		Express Mail No. EV 841 072 613 US

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm Name	BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, LLP 12400 Wilshire Boulevard, Seventh Floor, Los Angeles, California 90025-1030		
Signature			
Printed name	James C. Scheller		
Date	7/11/2006	Reg. No.	31,195

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: Mail Stop PCT, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date shown below:

Signature		Date	7/11/2006
Typed or printed name	Jessica Savage		

This collection of information is required by 37 CFR 1.5. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: MS PCT, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

中华人民共和国国家知识产权局
STATE INTELLECTUAL PROPERTY OFFICE
OF THE PEOPLE'S REPUBLIC OF CHINA



证 明 CERTIFICATE

本证明之附件是向中国专利局作为受理局提交的下列国际申请副本
IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY OF THE BELOW
IDENTIFIED INTERNATIONAL APPLICATION THAT WAS FILED WITH THE
CHINESE PATENT OFFICE AS RECEIVING OFFICE

申请号: PCT/CN2005/002304

NATIONAL APPLICATION NUMBER

申请日: 23. 12 月 2005 (23.12.2005)

NATIONAL FILING DATE

名称: MANAGING DEVICE MODELS IN A VIRTUAL MACHINING
INVENTION CLUSTER ENVIRONMENT

CERTIFIED COPY OF
PRIORITY DOCUMENT

中华人民共和国国家知识产权局局长
COMMISSIONER OF THE STATE INTELLECTUAL PROPERTY
OFFICE OF THE PEOPLE'S REPUBLIC OF CHINA

司力普

二零零六年五月十七日

MAY 17. 2006

文 理 本

PCT

REQUEST

The undersigned requests that the present international application be processed according to the Patent Cooperation Treaty.

For receiving Office use only

PCT/CN 2005 1002304
International Application No.

2 International Filing Date 23.12.2005

RO/CN 中华人民共和国国家知识产权局
PCT International Application

Name of receiving Office and "PCT International Application"

Applicant's or agent's file reference
(if desired) (12 characters maximum) FPEL05150068

Box No. I TITLE OF INVENTION
MANAGING DEVICE MODELS IN A VIRTUAL MACHINE CLUSTER ENVIRONMENT

Box No. II APPLICANT This person is also inventor

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

INTEL CORPORATION
2200 Mission College Blvd.
Santa Clara, California 95052
United States of America

Telephone No.

Facsimile No.

Teleprinter No.

Applicant's registration No. with the Office

State (that is, country) of nationality:
US

State (that is, country) of residence:
US

This person is applicant all designated States all designated States except the United States of America the United States of America only the States indicated in the Supplemental Box

Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

LIU, Benjamin
Room 302, No. 264 South Shanxi Road
Shanghai 200001
P. R. of China

This person is:

applicant only

applicant and inventor

inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:
CN

State (that is, country) of residence:
CN

This person is applicant all designated States all designated States except the United States of America the United States of America only the States indicated in the Supplemental Box

Further applicants and/or (further) inventors are indicated on a continuation sheet.

Box No. IV AGENT OR COMMON REPRESENTATIVE; OR ADDRESS FOR CORRESPONDENCE

The person identified below is hereby/has been appointed to act on behalf of the applicant(s) before the competent International Authorities as:

agent

common representative

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country.)

China Patent Agent (H.K.) Ltd.
22/F, Great Eagle Centre
23 Harbour Road, Wanchai
Hong Kong Special Administrative Region
The People's Republic of China

Telephone No.
(852)28284688

Facsimile No.
(852)28271018

Teleprinter No.

Agent's registration No. with the Office

Address for correspondence: Mark this check-box where no agent or common representative is/has been appointed and the space above is used instead to indicate a special address to which correspondence should be sent.

Sheet No. ... 2 ...

Continuation of Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)

If none of the following sub-boxes is used, this sheet should not be included in the request.

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

LING, Xiaofeng
Room 201, 2# Bldg,
No. 61 Haiting Road
Shanghai 201204
P. R. of China

This person is:

applicant only
 applicant and inventor
 inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

CN

State (that is, country) of residence:

CN

This person is applicant for the purposes of:

 all designated States all designated States except the United States of America the United States of America only the States indicated in the Supplemental Box

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

JIANG, Yunhong
Room 2107, No. 3 Lane 652
Daduhe Road
Shanghai 200001
P. R. of China

This person is:

applicant only
 applicant and inventor
 inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

CN

State (that is, country) of residence:

CN

This person is applicant for the purposes of:

 all designated States all designated States except the United States of America the United States of America only the States indicated in the Supplemental Box

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

LI, Xin
Room 302,
No. 264 South Shanxi Road
Shanghai 200001
P. R. of China

This person is:

applicant only
 applicant and inventor
 inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

CN

State (that is, country) of residence:

CN

This person is applicant for the purposes of:

 all designated States all designated States except the United States of America the United States of America only the States indicated in the Supplemental Box

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

This person is:

applicant only
 applicant and inventor
 inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

State (that is, country) of residence:

This person is applicant for the purposes of:

 all designated States all designated States except the United States of America the United States of America only the States indicated in the Supplemental Box Further applicants and/or (further) inventors are indicated on another continuation sheet.

Box No. V DESIGNATIONS

The filing of this request constitutes under Rule 4.9(a), the designation of all Contracting States bound by the PCT on the international filing date, for the grant of every kind of protection available and, where applicable, for the grant of both regional and national patents.

However,

DE Germany is not designated for any kind of national protection
 KR Republic of Korea is not designated for any kind of national protection
 RU Russian Federation is not designated for any kind of national protection

(The check-boxes above may be used to exclude (irrevocably) the designations concerned in order to avoid the ceasing of the effect, under the national law, of an earlier national application from which priority is claimed. See the Notes to Box No. V as to the consequences of such national law provisions in these and certain other States.)

Box No. VI PRIORITY CLAIM

The priority of the following earlier application(s) is hereby claimed:

Filing date of earlier application (day/month/year)	Number of earlier application	Where earlier application is:		
		national application: country or Member of WTO	regional application:*	international application: receiving Office
item (1)				
item (2)				
item (3)				

Further priority claims are indicated in the Supplemental Box.

The receiving Office is requested to prepare and transmit to the International Bureau a certified copy of the earlier application(s) (only if the earlier application was filed with the Office which for the purposes of this international application is the receiving Office) identified above as:

all items item (1) item (2) item (3) other, see Supplemental Box

* Where the earlier application is an ARIPO application, indicate at least one country party to the Paris Convention for the Protection of Industrial Property or one Member of the World Trade Organization for which that earlier application was filed (Rule 4.10(b)(ii)):

Box No. VII INTERNATIONAL SEARCHING AUTHORITY

Choice of International Searching Authority (ISA) (if two or more International Searching Authorities are competent to carry out the international search, indicate the Authority chosen; the two-letter code may be used):

ISA / CN

Request to use results of earlier search; reference to that search (if an earlier search has been carried out by or requested from the International Searching Authority):

Date (day/month/year) Number Country (or regional Office)

Box No. VIII DECLARATIONS

The following declarations are contained in Boxes Nos. VIII (i) to (v) (mark the applicable check-boxes below and indicate in the right column the number of each type of declaration):

Number of declarations

<input type="checkbox"/> Box No. VIII (i)	Declaration as to the identity of the inventor	:
<input type="checkbox"/> Box No. VIII (ii)	Declaration as to the applicant's entitlement, as at the international filing date, to apply for and be granted a patent	:
<input type="checkbox"/> Box No. VIII (iii)	Declaration as to the applicant's entitlement, as at the international filing date, to claim the priority of the earlier application	:
<input type="checkbox"/> Box No. VIII (iv)	Declaration of inventorship (only for the purposes of the designation of the United States of America)	:
<input type="checkbox"/> Box No. VIII (v)	Declaration as to non-prejudicial disclosures or exceptions to lack of novelty	:

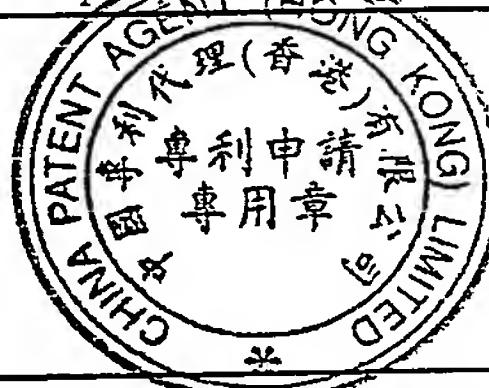
Sheet No. 4

Box No. IX CHECK LIST; LANGUAGE OF FILING

This international application contains:		This international application is accompanied by the following item(s) (mark the applicable check-boxes below and indicate in right column the number of each item):		Number of items
(a) in paper form, the following number of sheets:				
request (including declaration sheets)	: 4	<input checked="" type="checkbox"/> fee calculation sheet		: 1
description (excluding sequence listing and/or tables related thereto)	: 14	<input checked="" type="checkbox"/> original separate power of attorney		: 1
claims	: 11	<input type="checkbox"/> original general power of attorney		:
abstract	: 1	<input type="checkbox"/> copy of general power of attorney; reference number, if any:		:
drawings	: 6	<input type="checkbox"/> statement explaining lack of signature		:
Sub-total number of sheets	: 36	<input type="checkbox"/> priority document(s) identified in Box No. VI as item(s):		:
sequence listing	:	<input type="checkbox"/> translation of international application into (language):		:
tables related thereto	:	<input type="checkbox"/> separate indications concerning deposited microorganism or other biological material		:
(for both, actual number of sheets if filed in paper form, whether or not also filed in computer readable form; see (c) below)		<input type="checkbox"/> sequence listing in computer readable form (indicate type and number of carriers)		:
Total number of sheets	: 36	(i) <input type="checkbox"/> copy submitted for the purposes of international search under Rule 13ter only (and not as part of the international application)		:
(b) <input type="checkbox"/> only in computer readable form (Section 801(a)(i))		(ii) <input type="checkbox"/> (only where check-box (b)(i) or (c)(i) is marked in left column) additional copies including, where applicable, the copy for the purposes of international search under Rule 13ter		:
(i) <input type="checkbox"/> sequence listing		(iii) <input type="checkbox"/> together with relevant statement as to the identity of the copy or copies with the sequence listing mentioned in left column		:
(ii) <input type="checkbox"/> tables related thereto				
(c) <input type="checkbox"/> also in computer readable form (Section 801(a)(ii))				
(i) <input type="checkbox"/> sequence listing				
(ii) <input type="checkbox"/> tables related thereto				
Type and number of carriers (diskette, CD-ROM, CD-R or other) on which are contained the				
<input type="checkbox"/> sequence listing:				
<input type="checkbox"/> tables related thereto:				
(additional copies to be indicated under items 9(ii) and/or 10(ii), in right column)				
Figure of the drawings which should accompany the abstract:	Language of filing of the international application: EN			

Box No. X SIGNATURE OF APPLICANT, AGENT OR COMMON REPRESENTATIVE

Next to each signature, indicate the name of the person signing and the capacity in which the person signs (if such capacity is not obvious from reading the request).



For receiving Office use only

1. Date of actual receipt of the purported international application:

23 · 12月 2005 (23 · 12 · 2005)

2. Drawings:

 received:

3. Corrected date of actual receipt due to later but timely received papers or drawings completing the purported international application:

4. Date of timely receipt of the required corrections under PCT Article 11(2):

5. International Searching Authority (if two or more are competent): ISA /

6. Transmittal of search copy delayed until search fee is paid not received:

For International Bureau use only

Date of receipt of the record copy by the International Bureau:

See Notes to the request form

MANAGING DEVICE MODELS IN A VIRTUAL MACHINE CLUSTER ENVIRONMENT

BACKGROUND

[0001] A virtual machine (VM) architecture logically partitions a physical machine, 5 such that the underlying hardware of the machine is time-shared and appears as one or more independently operation virtual machines. A computer platform in a virtual machine cluster environment may comprise a virtual machine monitor (VMM) that may create a plurality of virtual machines and runs on the computer platform to facilitate for other software the abstraction of one or more 10 virtual machines. The computer platform may further comprise a plurality of device models that may be virtualizations/simulations of real devices. The virtual machine monitor may manage resource usage for operations between the virtual machines and the device models, such as device related input/output operation and device model initiated interrupt. Examples for the 15 virtual machine monitor may comprise a hybrid virtual machine monitor, a host virtual machine monitor and a hypervisor virtual machine monitor. Examples for the real device may comprise various input/output (I/O) devices, an interrupt controller, and an event timer, etc.

BRIEF DESCRIPTION OF THE DRAWINGS

20[0002] The invention described herein is illustrated by way of example and not by way of limitation in the accompanying figures. For simplicity and clarity of illustration, elements illustrated in the figures are not necessarily drawn to scale. For example, the dimensions of some elements may be exaggerated relative to other elements for clarity. Further, where considered appropriate,

reference labels have been repeated among the figures to indicate corresponding or analogous elements.

- [0003] Fig. 1 shows an embodiment of a virtual machine cluster environment.
- [0004] Fig. 2 shows an embodiment of structures of a server platform and a plurality of client platforms in the virtual machine cluster environment of Fig. 1.
- 5 [0005] Fig. 3 shows an embodiment of a method of implementing a device related input/output operation in the virtual machine cluster environment of Fig. 2.
- [0006] Fig. 4 shows an embodiment of a method of implementing a device model initiated interrupt in the virtual machine cluster environment of Fig. 2.
- 10 [0007] Fig. 5 shows an embodiment of a computer platform that may be implemented as a platform within the virtual machine cluster environment of Fig. 2.
- [0008] Fig. 6 shows another embodiment of structures of a server platform and a plurality of client platforms in the virtual machine cluster environment of Fig. 1.

15 DETAILED DESCRIPTION

- [0009] The following description describes techniques for managing device models in a virtual machine cluster environment. In the following description, numerous specific details such as logic implementations, pseudo-code, means to specify operands, resource partitioning/sharing/duplication implementations, 20 types and interrelationships of system components, and logic partitioning/integration choices are set forth in order to provide a more thorough understanding of the current invention. However, the invention may be practiced without such specific details. In other instances, control structures, gate level circuits and full software instruction sequences have not

been shown in detail in order not to obscure the invention. Those of ordinary skill in the art, with the included descriptions, will be able to implement appropriate functionality without undue experimentation.

[0010] References in the specification to "one embodiment", "an embodiment", "an example embodiment", etc., indicate that the embodiment described may 5 include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is 10 described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

[0011] Embodiments of the invention may be implemented in hardware, firmware, 15 software, or any combination thereof. Embodiments of the invention may also be implemented as instructions stored on a machine-readable medium, that may be read and executed by one or more processors. A machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computing device). For example, a 20 machine-readable medium may include read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.) and others.

[0012] An embodiment of a virtual machine cluster environment 1 is shown in Fig.

1. The virtual machine cluster environment 1 may comprise a plurality of client platforms 10/30 and a server platform 20 that may connect with the client platforms 10/30 through a network 40, such as Ethernet, fiber channel or other like communications link. In the embodiment, the virtual machine cluster environment 1 may comprise any number of client platforms and server platforms.

[0013] Each of the client platforms 10/30 may host a plurality of virtual machines running their own guest operating system(s) and guest application software.

10 The server platform 20 may provide a variety of services to the virtual machines of the client platforms, for example, but not limited to, device virtualization/simulation services and virtual machine management services. A non-exhaustive list of examples for the client platform 10/30 may include mainframe computer, mini-computer, personal computer, portable computer, 15 laptop computer and other devices for transceiving and processing data. A non-exhaustive list of examples for the server platform 20 may include supercomputer, mainframe computer, workstation and other devices for transceiving and processing data.

[0014] An embodiment of structures for the client platforms 10/30 and the server 20 in the virtual machine cluster environment 1 is shown in Fig. 2. The client platforms 10/30 and server platform 20 may be configured according to the hybrid virtual machine monitor architecture.

[0015] In the embodiment, each of the client platforms 10/30 may comprise an underlying hardware machine 11/31 having one or more processors, memory, 25 I/O devices, etc., wherein the I/O devices may further comprise a network



interface to transceive data through the network 40, such as the network interface 110 of the client platform 10 and the network interface 310 of the client platform 30.

[0016] Each of the client platforms 10/30 may further comprise a hypervisor (or 5 micro-hypervisor) 12/32, which is a kernel component configured according to the hybrid virtual machine monitor architecture and responsible for processor(s)/memory virtualization, interrupt handling and domain scheduling. The hypervisor 12/32 may further comprise a network interface management core 120/320 to manage the network interface 110/310 of the underlying 10 machine. Examples for the hypervisor may comprise Xen's hypervisor, Virtual PC's hypervisor, and the like.

[0017] Each of the client platforms 10/30 may further comprise a plurality of guest 15 software images running on a plurality of virtual machines created and managed by the hypervisor 12/32, for example, guest operating systems 131/141 and guest applications 132/142 running on the virtual machines 13/14 of the client platform 10, and guest operating systems 331/341 and guest applications 332/342 running on the virtual machines 33/34 of the client 20 platform 30. The guest operating systems 131, 141, 331 and 341 may be different with each other. The number of guest operating systems running on the client platform may depend on utilization of processor/memory resources and device resources. For example, the number of guest operating systems may be low if the client platform utilizes the processor/memory resources more intensively compared to device utilization. However, the number of guest 25 operating systems may be high if the client platform utilizes the device resources more intensively compared to processor/memory utilization.

Examples for the guest operating system may comprise Linux, Windows, FreeBSD Unix, etc.

[0018] The server platform 20 may comprise an underlying hardware machine 21 having one or more processors, memory, I/O devices, etc., wherein the I/O devices 5 may further comprise a network interface 210 to transceive data through the network 40. The server platform 20 may be loaded with a service operating system 22 configured according to the hybrid virtual machine monitor architecture. The service operating system 22 may respectively interact with the hypervisors 12/32 of the client platforms 10/30 via the network 10 40 to service the plurality of virtual machines in the client platforms 10/30.

Examples for the service operating system may comprise Linux, Windows, FreeBSD Unix, etc.

[0019] The server platform 20 may be further loaded with a plurality of device models 23 and a control panel 24. The device models 23 may be virtual devices 15 that may be created or defined for example according to the hybrid virtual machine monitor architecture. In some embodiments, because of device virtualization/simulation nature, the device models 23 may not be connected to or represented by a real instance of a device, and may not be reflected in a real device that is connected to a hardware component. Examples for the 20 device models 23 may comprise, but not limited to, various I/O devices, an interrupt controller, an event timer, etc.

[0020] The device models 23 may work with the service operating system 22 to provide device virtualization/simulation service to the virtual machines of the client platforms 10/30 via the network 40, including hosting device related 25 operation from/to the virtual machines, such as device related I/O operation

and device model initiated interrupt. In some embodiments, the device models in the virtual machine cluster environment may be centralized in the server platform 20. However, in other embodiments, the client platform 10/30 may keep some of the device models, such as the device models frequently used by the virtual machines of the client platform. For example, the client platform 10/30 may keep the device models inside of the hypervisor 11/31.

5 [0021] The control panel 24 may work with the service operating system 22 to manage the virtual machines of the client platforms 10/30 via the network 40. In an embodiment, the control panel 24 may comprise a database 241 containing identifiers of the virtual machines (or the guest operating systems running on the virtual machines), the hypervisors and the client platforms in the virtual machine cluster environment that are served by the server platform 20; and as another possibility the database can be created or updated on fly. In an embodiment, a virtual machine of the client platform 10/30 may send a request 10 for a device related input/output operation to the server platform 20 via the network 40. The control panel 24 and the service operating system 22 may identify a corresponding device model from the plurality of device models 23 to host the request, and further construct a feedback packet by incorporating a result of the input/output operation from the device model with identifier(s) to 15 identify the virtual machine that has sent the operation request. In another embodiment, a device model may initiate an interrupt for a virtual machine of the client platform 10/30. The control panel 24 and the service operating system 22 may construct an interrupt packet by incorporating the interrupt instruction with identifier(s) to identify the virtual machine that may handle the 20 interrupt. The identifier(s) may comprise a virtual machine identifier, a 25

hypervisor identifier and/or client platform identifier. In addition, the control panel 24 may perform virtual machine management for other aspect, such as virtual machine operation management including virtual machine launching management, virtual machine shutdown management, virtual machine migration management, virtual machine suspend management, etc.

5 [0022] The server platform 20 may further be loaded with a plurality of application software running on the service operating system 22, such as applications 25 and 26.

10 [0023] Other embodiments may implement other technologies to the structure depicted in Fig. 2. In an embodiment, the server platform 20 may further comprise a hypervisor running between the underlying hardware machine 21 and the service operating system 22. In such case, the server platform 20 may further comprise one or more virtual machine(s) defined by the hypervisor and the service operating system. In another embodiment, console devices (e.g., 15 frequently used input/output devices such as a keyboard, a mouse, a video device, etc.) may be centralized in the underlying hardware machine 21 of the server platform 20. In such case, the client platform 10/30 may not be provided with the console devices.

20 [0024] Fig. 3 shows an embodiment of a method of implementing a device related input/output operation in the virtual machine cluster environment 1.

25 [0025] In the embodiment, an unauthorized input/output operation for inputting a data from a device (input operation) or outputting a data to the device (output operation) happens in a guest operating system running on a virtual machine of a client platform (e.g., guest operating system 131 running on the virtual machine 13 of the client platform 10), and a corresponding device driver in the

guest operating system may execute an 'IN' instruction for input operation or 'OUT' instruction for output operation, which may cause a trap into the hypervisor of the client platform in block 301, because the guest operating system is deprivileged.

5[0026] In block 302, the hypervisor of the client platform may perceive the unauthorized I/O operation happened in the guest operating system through a virtual machine exit, which is a transition from non-root VMX operation in the virtual machine to root VMX operation in the hypervisor.

[0027] In block 303, the hypervisor of the client platform may construct an I/O request packet and send the packet to the server platform 20 through the network 40. In an embodiment, the I/O request may further comprise identifier(s) to identify a device model in the server platform 20 corresponding to the device that is related to the I/O operation in the guest operating system of the virtual machine. The identifier(s) may comprise a device model identifier and/or a server platform identifier (e.g., a network ID for the server platform). In response to receiving the I/O request packet, the service operating system 22 may identify the corresponding device model from the plurality of device models 23 under control from the control panel 24 and route the I/O request packet to the identified device model in block 304.

20[0028] In block 305, the identified device model may handle the I/O request. For output request, the data in the I/O request packet may be output to the identified device model. Depending upon inter-platform communication interface definition, the server platform 20 may send back a finish notification to the I/O requester (i.e., the client platform 10/30) to notify finish of the 'OUT' operation. However, for input request, the identified device model 23 may

obtain a feedback to the input request, for example, through cooperating with the service operating system 22 and the underlying hardware 21 of the service platform 20, and send a feedback packet to the service operating system 22 in block 306. The feedback packet may incorporate the feedback with identifier(s) to identify the guest operating system (or virtual machine) of the client platform that is executing the I/O operation. In an embodiment, the identifier(s) may comprise a guest operating system identifier (or virtual machine identifier) and/or a client platform identifier (e.g., a network ID for the client platform). In another embodiment, if the client platforms in the virtual machine cluster environment 1 comprises different types of hypervisors, e.g., Xen hypervisor, Virtual PC's hypervisor, etc., then the identifier may further comprise a hypervisor identifier to identify the different hypervisors.

[0029] In block 307, under control from the control panel stored with identifiers of the guest operating systems (or virtual machines), the hypervisors and the client platforms in the virtual machine cluster environment 1, the service operating system 22 may route the feedback packet to the client platform that has sent the I/O request based upon the identifier(s) in the feedback packet. In block 308, the hypervisor of the client platform may identify the guest operating system (or virtual machine) executing the corresponding I/O operation and send the feedback packet to the identified guest operating system through a virtual machine entry, which is another transition from the root VMX operation in the hypervisor to the non-root VMX operation in the virtual machine.

[0030] An embodiment of a method of implementing a device model initiated interrupt in the virtual machine cluster environment 1 is shown in Fig. 4.

[0031] In block 401, one of the device models 23 in the server platform 20 initiate an interrupt for a guest operating system of a client platform in the virtual machine cluster environment 1. In block 402, the device model passes an interrupt packet incorporating the interrupt instruction as well as identifier(s) to identify the guest operating system (or virtual machine) of the client platform to the service operating system 22. In an embodiment, the identifier(s) may comprise a guest operating system identifier (or virtual machine identifier) and/or a client platform identifier (e.g., a network ID for the client platform). In another embodiment, if the client platforms in the virtual machine cluster environment 1 comprises different types of hypervisors, e.g., Xen hypervisor, Virtual PC's hypervisor, etc., then the identifier may further comprise a hypervisor identifier to identify the different hypervisors.

[0032] In block 403, under control from the control panel 24, the service operating system 22 may send the interrupt packet to the client platform identified by the identifier(s) in the interrupt packet. The hypervisor of the client platform may further identify the guest operating system based upon the interrupt packet and inject the interrupt to the identified guest operating system through the virtual machine entry (block 404), so that the guest operating system may handle the interrupt (block 405).

[0033] Fig. 5 shows an embodiment of a computer platform that may be implemented as a platform in the virtual machine cluster environment 1.

[0034] In an embodiment, the computing platform may comprise one or more processors 50, memory 51, chipset 52, I/O device 53, BIOS firmware 54 and the like. The one or more processors 50 are communicatively coupled to various components (e.g., the memory 51) via one or more buses such as a

processor bus. The processors 50 may be implemented as an integrated circuit (IC) with one or more processing cores that may execute codes under a suitable architecture, for example, including Intel® Xeon™, Intel® Pentium™, Intel® Itanium™ architectures, available from Intel Corporation of Santa Clara, California.

5

[0035] In an embodiment, the memory 51 may store codes to be executed by the processor 50. A non-exhaustive list of examples for the memory 51 may comprise one or a combination of the following semiconductor devices, such as synchronous dynamic random access memory (SDRAM) devices, 10 RAMBUS dynamic random access memory (RDRAM) devices, double data rate (DDR) memory devices, static random access memory (SRAM), flash memory devices, and the like.

10

[0036] In an embodiment, the chipset 52 may provide one or more communicative path among the processor 50, memory 51 and various components, such as 15 the I/O device 53, peripheral component 54, and BIOS firmware 55. The chipset 52 may comprise a memory controller hub 520, an input/output controller hub 521 and a firmware hub 522.

15

20

[0037] In an embodiment, the memory controller hub 520 may provide a communication link to the processor bus that may connect with the processor 20 50 and to a suitable device such as the memory 51. The memory controller hub 520 may couple with the I/O controller hub 521 that may provide an interface to the I/O devices 53. A non-exhaustive list of examples for the I/O devices 53 may comprise a keyboard, a mouse, a storage device, a video device, an audio device, a network interface, and the like.

[0038] In an embodiment, the memory controller hub 520 may communicatively couple with a firmware hub 522 via the input/output controller hub 521. The firmware hub 522 may couple with the BIOS firmware 54 that may store routines that the computing platform executes during system startup in order to 5 initialize the processors 50, chipset 52, and other components of the computing platform. Moreover, the BIOS firmware 54 may comprise routines or drivers that the computing platform may execute to communicate with one or more components of the compute platform.

[0039] In an embodiment, the computer platform as depicted in Fig. 5 may perform 10 as a client 10/30. The memory 51 may store software images such as a hypervisor, guest operating systems, guest applications, etc. In another embodiment, the computer platform as depicted in Fig. 5 may perform as a server platform 20. The memory 51 may store software images such as a service operating system, device models, control panel, and optionally various 15 applications, etc.

[0040] Another embodiment for the server platform and client platform structures in the virtual machine cluster environment 1 is depicted in Fig. 6. As depicted, more than one server platforms (e.g., owner server platform 620 and I/O server platform 6200) may connect with the client platform 10/30 through the 20 network 40 to provide them with device virtualization/simulation services. The owner server platform 620 has similar structure as the server platform 20 and the I/O server platform 6200 is loaded with a plurality of device models 6230 so that the server platforms 620/6200 may provide robust device virtualization 25 services to their client platforms. Functionalities of the control panel 624 may encompass management of the device models 6230 in the I/O server platform

6200. In an embodiment, the control panel 624 may perform the device model identification based upon an I/O request packet from the client platform 10/30. In such case, the owner server platform 620 may first receive the I/O request packet from the client platform 10/30 via the network 40 and route the I/O request packet to the device model 6230 in the I/O service platform 6200 after 5 analyzing the identifier(s) in the I/O request packet. In another embodiment, the device model 6230 may construct an I/O feedback packet or an interrupt packet by incorporating an identifier(s) to identifier the guest operating system (or virtual machine) in the intended client platform 10/30 under supervision 10 from the control panel 624.

[0041] While certain features of the invention have been described with reference to example embodiments, the description is not intended to be construed in a limiting sense. Various modifications of the example embodiments, as well as other embodiments of the invention, which are apparent to persons skilled in the art to which the invention pertains are deemed to lie within the spirit and scope of the invention. For example, although the above embodiments are described according to the hybrid virtual machine monitor architecture, the present invention may be applied to other kinds of virtual machine monitor architectures, such as a host virtual machine monitor. 15

What is claimed is:

1. A method for a client platform coupled to a server platform via a network, comprising:
 - 5 determining that an input/output operation related to a hardware device happens in a virtual machine of the client platform; and requesting the server platform via the network to handle the input/output operation related to the hardware device.
 - 10 2. The method of claim 1, wherein the request comprises a server identifier to identify the server platform.
 - 15 3. The method of claim 1, wherein the request comprises a software device identifier to identify a software device from a plurality of software devices in the server platform to handle the input/output operation related to the hardware device.
 - 20 4. The method of claim 1, further comprising: receiving a feedback for the input operation from the server through the network, the feedback comprising a virtual machine identifier to identify the virtual machine in the client platform that is executing the input operation; and sending the feedback to the virtual machine identified by the virtual machine identifier.

5. The method of claim 1, further comprising:

receiving via the network an interrupt instruction issued by a software device of the server platform, the interrupt instruction comprising a virtual machine identifier to identify a virtual machine to perform the interrupt instruction; and

5 injecting the interrupt instruction into the virtual machine identified by the virtual machine identifier.

6. A machine-readable medium comprising a plurality of instructions which when executed result in a client platform:

10 determining that an input/output operation related to a hardware device happens in a virtual machine of the client platform; and

 requesting a server platform coupled to the client platform via the network to handle the input/output operation related to the hardware device.

15 7. The machine-readable medium of claim 6, wherein the request further comprise a server identifier to identify the server platform.

8. The machine-readable medium of claim 6, wherein the request comprises a software device identifier to identify a software device from a plurality 20 of software devices in the server platform to handle the input/output operation related to the hardware device.

2/10

9. The machine-readable medium of claim 6, wherein the plurality of instructions further result in the client platform:

receiving a feedback for the input operation from the server platform through the network, the feedback comprising a virtual machine identifier to identify the virtual machine in the client platform that is executing the input operation; and

sending the feedback to the virtual machine identified by the virtual machine identifier.

10 10. The machine-readable medium of claim 6, wherein the plurality of instructions further result in the client platform:

receiving an interrupt instruction issued by a software device from the plurality of software devices in the server platform through the network, the interrupt instruction comprising a virtual machine identifier to identify a virtual machine to perform the interrupt instruction;

injecting the interrupt instruction to the virtual machine identified by the virtual machine identifier.

11. A method for a server platform coupled to a client platform via a network,

receiving, from the client platform via the network, a request for an input/output operation related to a hardware device; and

identifying a software device from a plurality of software devices in the server platform to handle the request, the identified software device corresponding to the hardware device related to the input/output operation.

12. The method of claim 11, wherein the request comprises a software device identifier to identify the software device in the server platform.

13. The method of claim 11, further comprising
5 receiving the request for the input/output operation with a first server of the server platform,
determining whether the identified software device is in a second server of the server platform; and
sending the request from the first server to the second server via the
10 network, in response to determining that the identified software device is in the second server.

14. The method of claim 11, further comprising:
obtaining a result for the input operation from the identified software device;
15 constructing a feedback with the result and a virtual machine identifier to identify a virtual machine in the client platform that is executing the input operation; and
sending the feedback from the server platform to the client platform through the network.
20

15. The method of claim 14, wherein the feedback further comprise a client platform identifier to identify the client platform that has sent the request.

16. The method of claim 11, further comprising:
25 issuing an interrupt instruction from a software device of the plurality of software devices in the server platform to the client platform through the network.



17. The method of claim 11, wherein the interrupt instruction further comprises a virtual machine identifier to identify a virtual machine in the client platform to handle the interrupt.

5 18. A machine-readable medium comprising a plurality of instructions which when executed result in a server platform:
 receiving a request for an input/output operation related to a hardware device from a client platform through a network; and
 identifying a software device from a plurality of software devices in the
10 server platform to handle the request, the identified software device corresponding to the hardware device related to the input/output operation.

19. The machine readable medium of claim 18, wherein the request comprises a software device identifier to identify the software device in the server
15 platform.

20. The machine readable medium of claim 18, wherein the plurality of instructions that result in the server platform receiving the request, further result in the server platform:
20 receiving the request for the input/output operation from the client platform with a first server in the server platform.

21. The machine readable medium of claim 20, wherein the plurality of instructions further result in the server platform:

determining whether the identified software device is in a second server of the server platform; and

5 sending the request from the first server to the second server through the network, in response to determining that the identified software device is in the second server.

22. The machine readable medium of claim 18, wherein the plurality of
10 instructions further result in the server platform:

obtaining a result for the input operation from the identified software device;

constructing a feedback with the result and a virtual machine identifier to identify a virtual machine in the client platform that is executing the input operation; and

15 sending the feedback from the server platform to the client platform through the network.

23. The machine readable medium of claim 22, wherein the feedback further comprise a client identifier to identify the client platform that has sent the
20 request.

24. The machine readable medium of claim 18, wherein the plurality of instructions further result in the server platform:

issuing an interrupt instruction from a software device of the plurality of
25 software devices in the server platform to the client platform through the network.

25. The machine readable medium of claim 24, wherein the interrupt instruction further comprises a virtual machine identifier to identify a virtual machine in the client platform to handle the interrupt.

5 26. A system, comprising
 a client platform comprising:
 a plurality of virtual machines; and
 a virtual machine monitor to determine that an input/output operation related to a hardware device happens in a virtual machine of the plurality of virtual
10 machines and construct a request for the input/output operation;
 a client network interface to send the request through a network; and
 the server platform comprising:
 a server network interface to receive the request through the network;
 a plurality of software devices;
15 a controller to identify a software device from the plurality of software devices to handle the request, the identified software device corresponding to the hardware device related to the input/output operation.

27. The system of claim 26, wherein the request further comprises a
20 software identifier to identifier the software device in the server platform.

28. The system of claim 26, wherein
the identified software device in the server platform is further to obtain a
result for the input operation, and construct a feedback with the result and a virtual
machine identifier to identify the virtual machine in the client platform under control
5 from the controller, and
the server network interface is further to send the feedback to the client
platform through the network.

29. The system of claim 26, wherein
10 the client network interface is further to receive a feedback for the input
operation from the server platform through the network; and
the virtual machine monitor is further to identify the virtual machine in the
client platform that is executing the input operation based upon the feedback and
send the feedback to the identified virtual machine.

15 30. The system of claim 26, wherein
a software device in the server platform is to issue an interrupt instruction
under control from the controller, the interrupt instruction including a virtual
machine identifier to identify another virtual machine in the client platform to
20 handle the interrupt instruction; and
the server network interface is further to send the interrupt instruction to the
client platform through the network.

31. The system of claim 30, wherein
the client network interface is further to receive the interrupt instruction; and
the virtual machine monitor is further to identify the another virtual machine
from the plurality of virtual machines based upon the interrupt instruction and
5 inject the interrupt into the identified another virtual machine.

32. A method for a system comprising a client platform and a server
platform, wherein the client platform couples to the server platform through a
network, the method comprising:

10 determining that an input/output operation related to a hardware device
happens in a virtual machine of the client platform;
sending a request for the input/output operation from the client platform to
the server platform through the network;
receiving the request through the network by the server platform;
15 identifying a software device from a plurality of software devices in the
server platform to handle the request, wherein the identified virtual device
corresponds to the hardware device related to the input/output operation.

33. The method of claim 32, wherein the receiving further comprises
20 receiving the request with a first server in the server platform.

34. The method of claim 33, further comprising:
determining whether the identified software device is in a second server of
the server platform; and
sending the request from the first server to the second server through the
5 network, in response to determining that the identified software device is in the
second server.

35. The method of claim 32, further comprising:
obtaining a result for the input operation from the identified software device
10 in the server platform;
constructing a feedback with the result and a virtual machine identifier to
identify the virtual machine in the client platform that is executing the input
operation; and
sending the feedback from the server platform to the client platform through
15 the network.

36. The method of claim 35, wherein the feedback further comprise a client
identifier to identify the client platform that has sent the request.

20 37. The method of claim 32, further comprising:
issuing an interrupt instruction from a software device in the server platform
to the client platform through the network.

38. The method of claim 32, wherein the interrupt instruction further
25 comprises a virtual machine identifier to identify another virtual machine in the
client platform to handle the interrupt.

39. The method of claim 38, further comprising:
receiving the interrupt instruction through the network by the client platform;
identifying the another virtual machine in the client platform based upon the
interrupt instruction; and
injecting the interrupt into the identified another virtual machine.



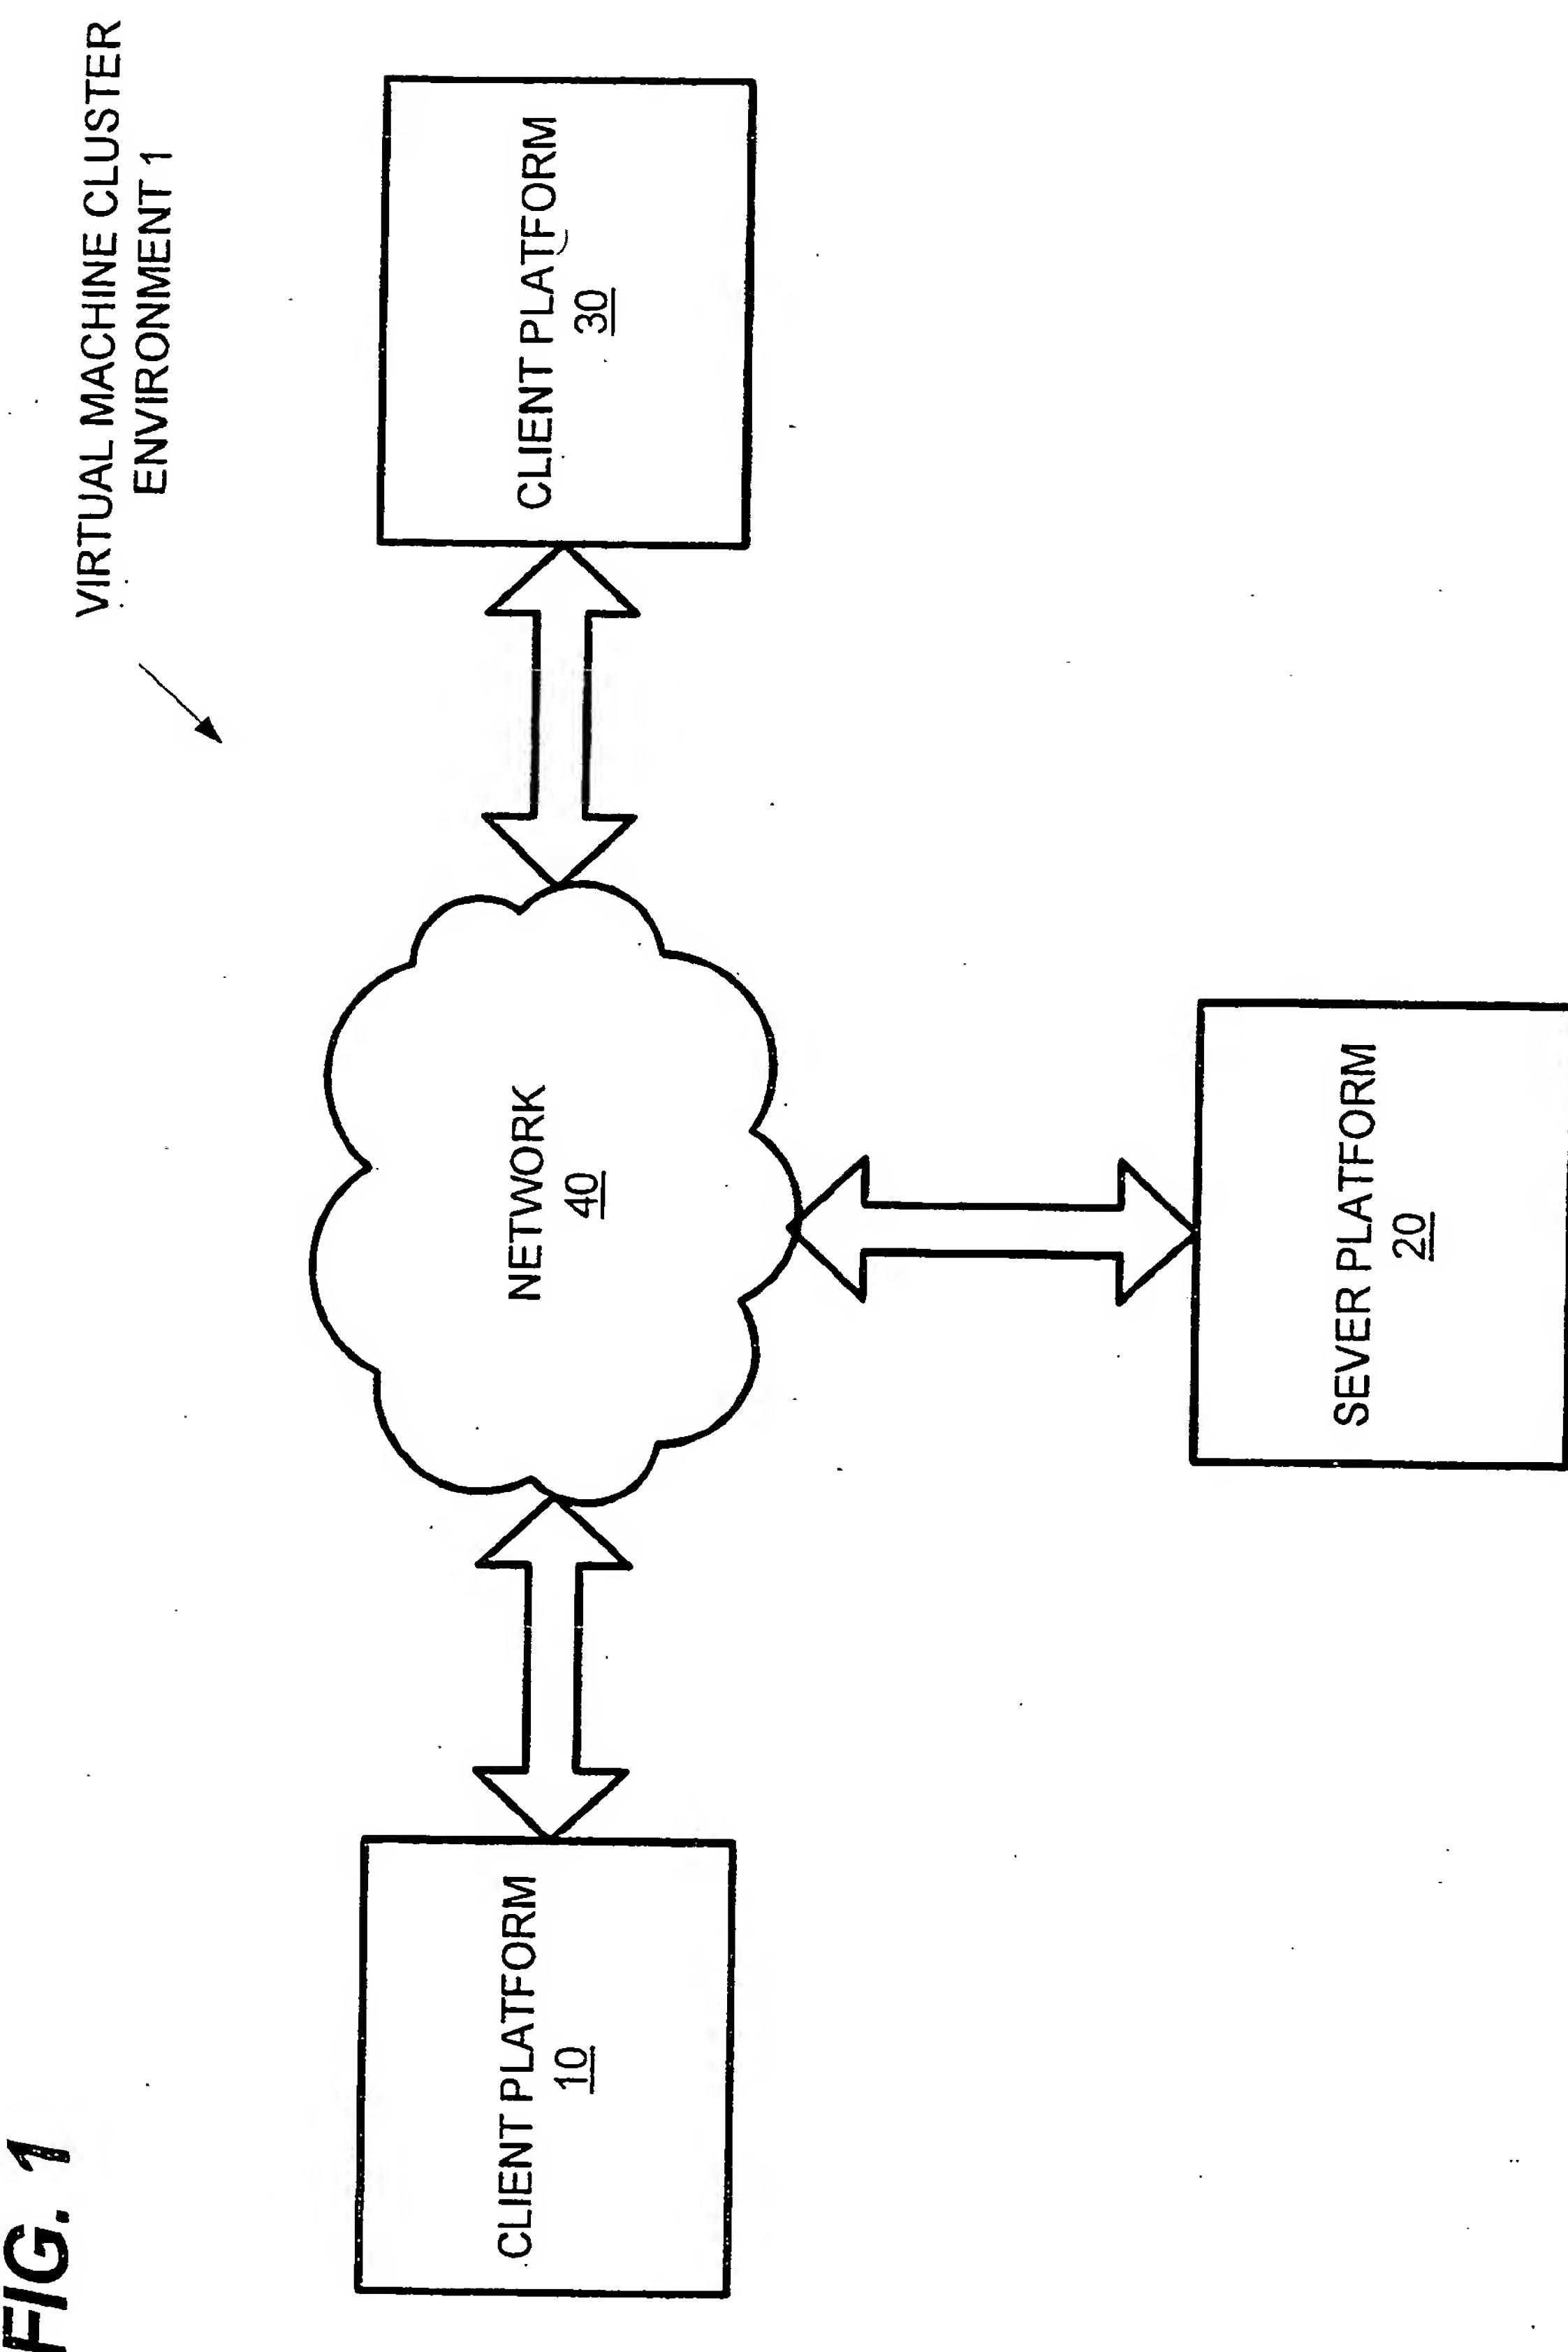
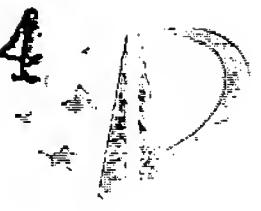
ABSTRACT

Machine-readable media, methods, apparatus and system are described.

In some embodiments, a client platform may determine that an input/output operation related to a hardware device may happen in a virtual machine of the

5 client platform and request a server platform via the network to handle the input/output operation related to the hardware. The server platform may receive the request from the client platform through the network; and identify a software device from a plurality of software devices in the server platform to handle the request, wherein the identified software device corresponds to the hardware

10 device related to the input/output operation.



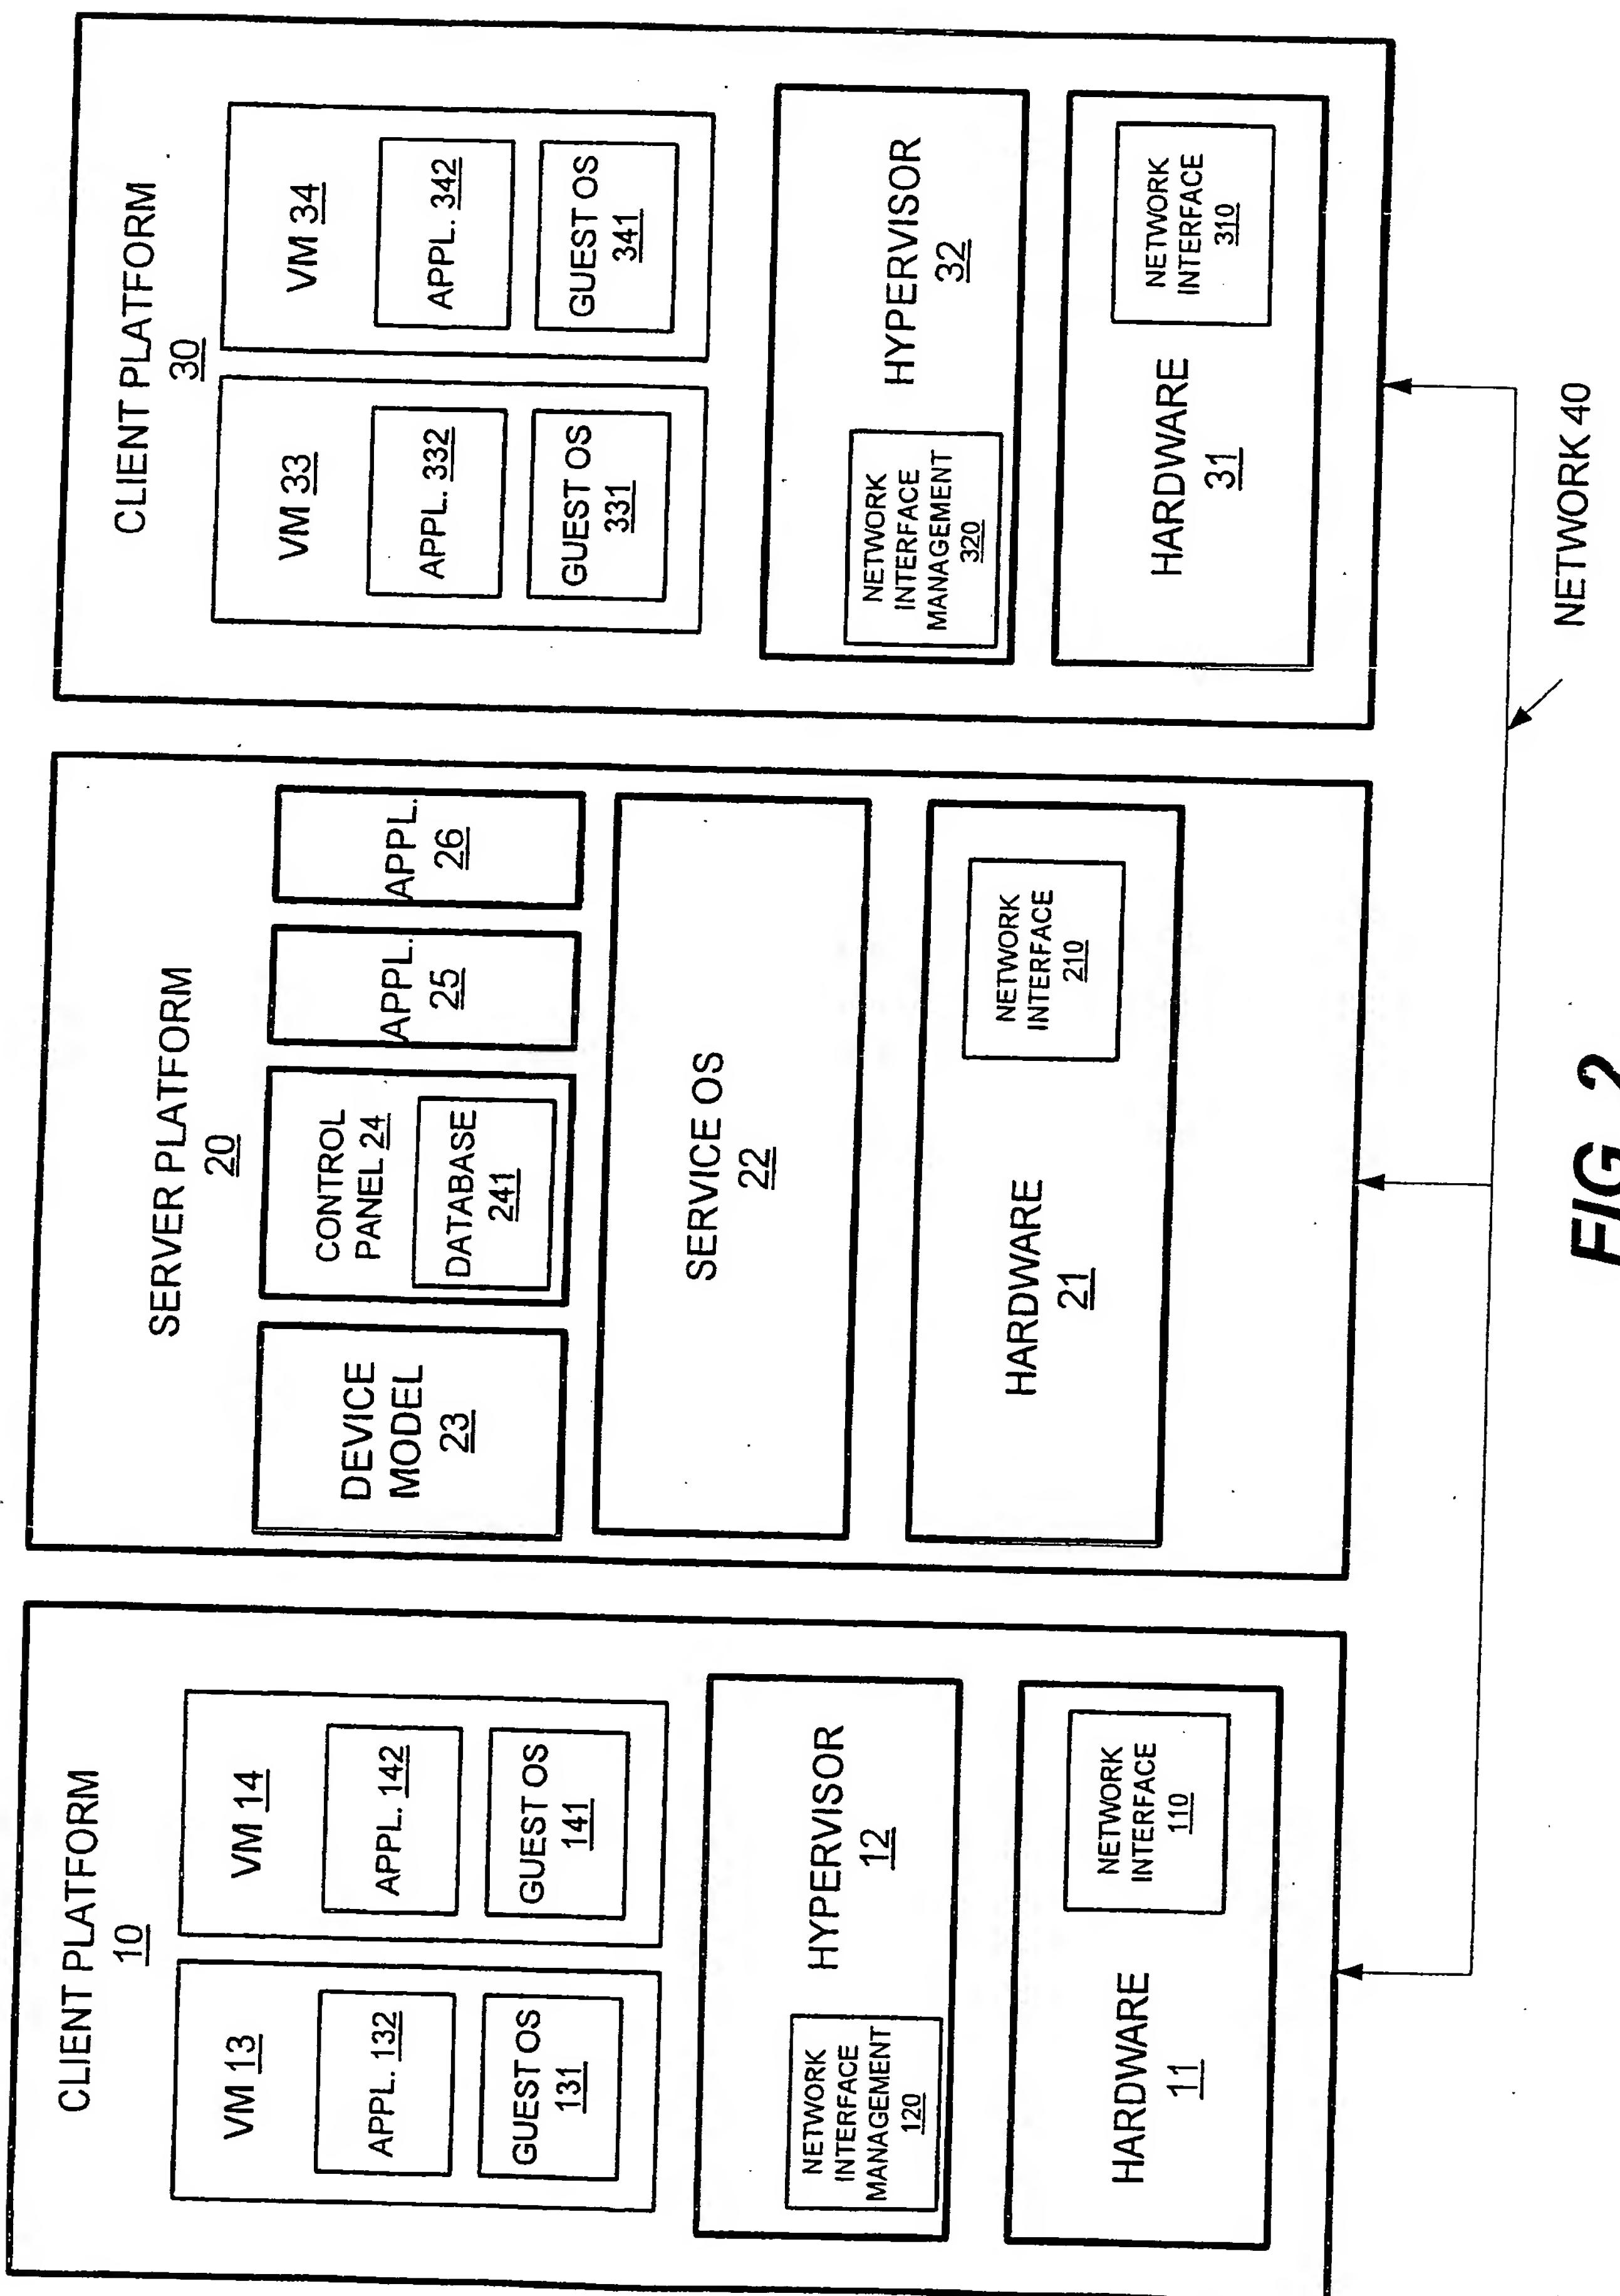


FIG. 2

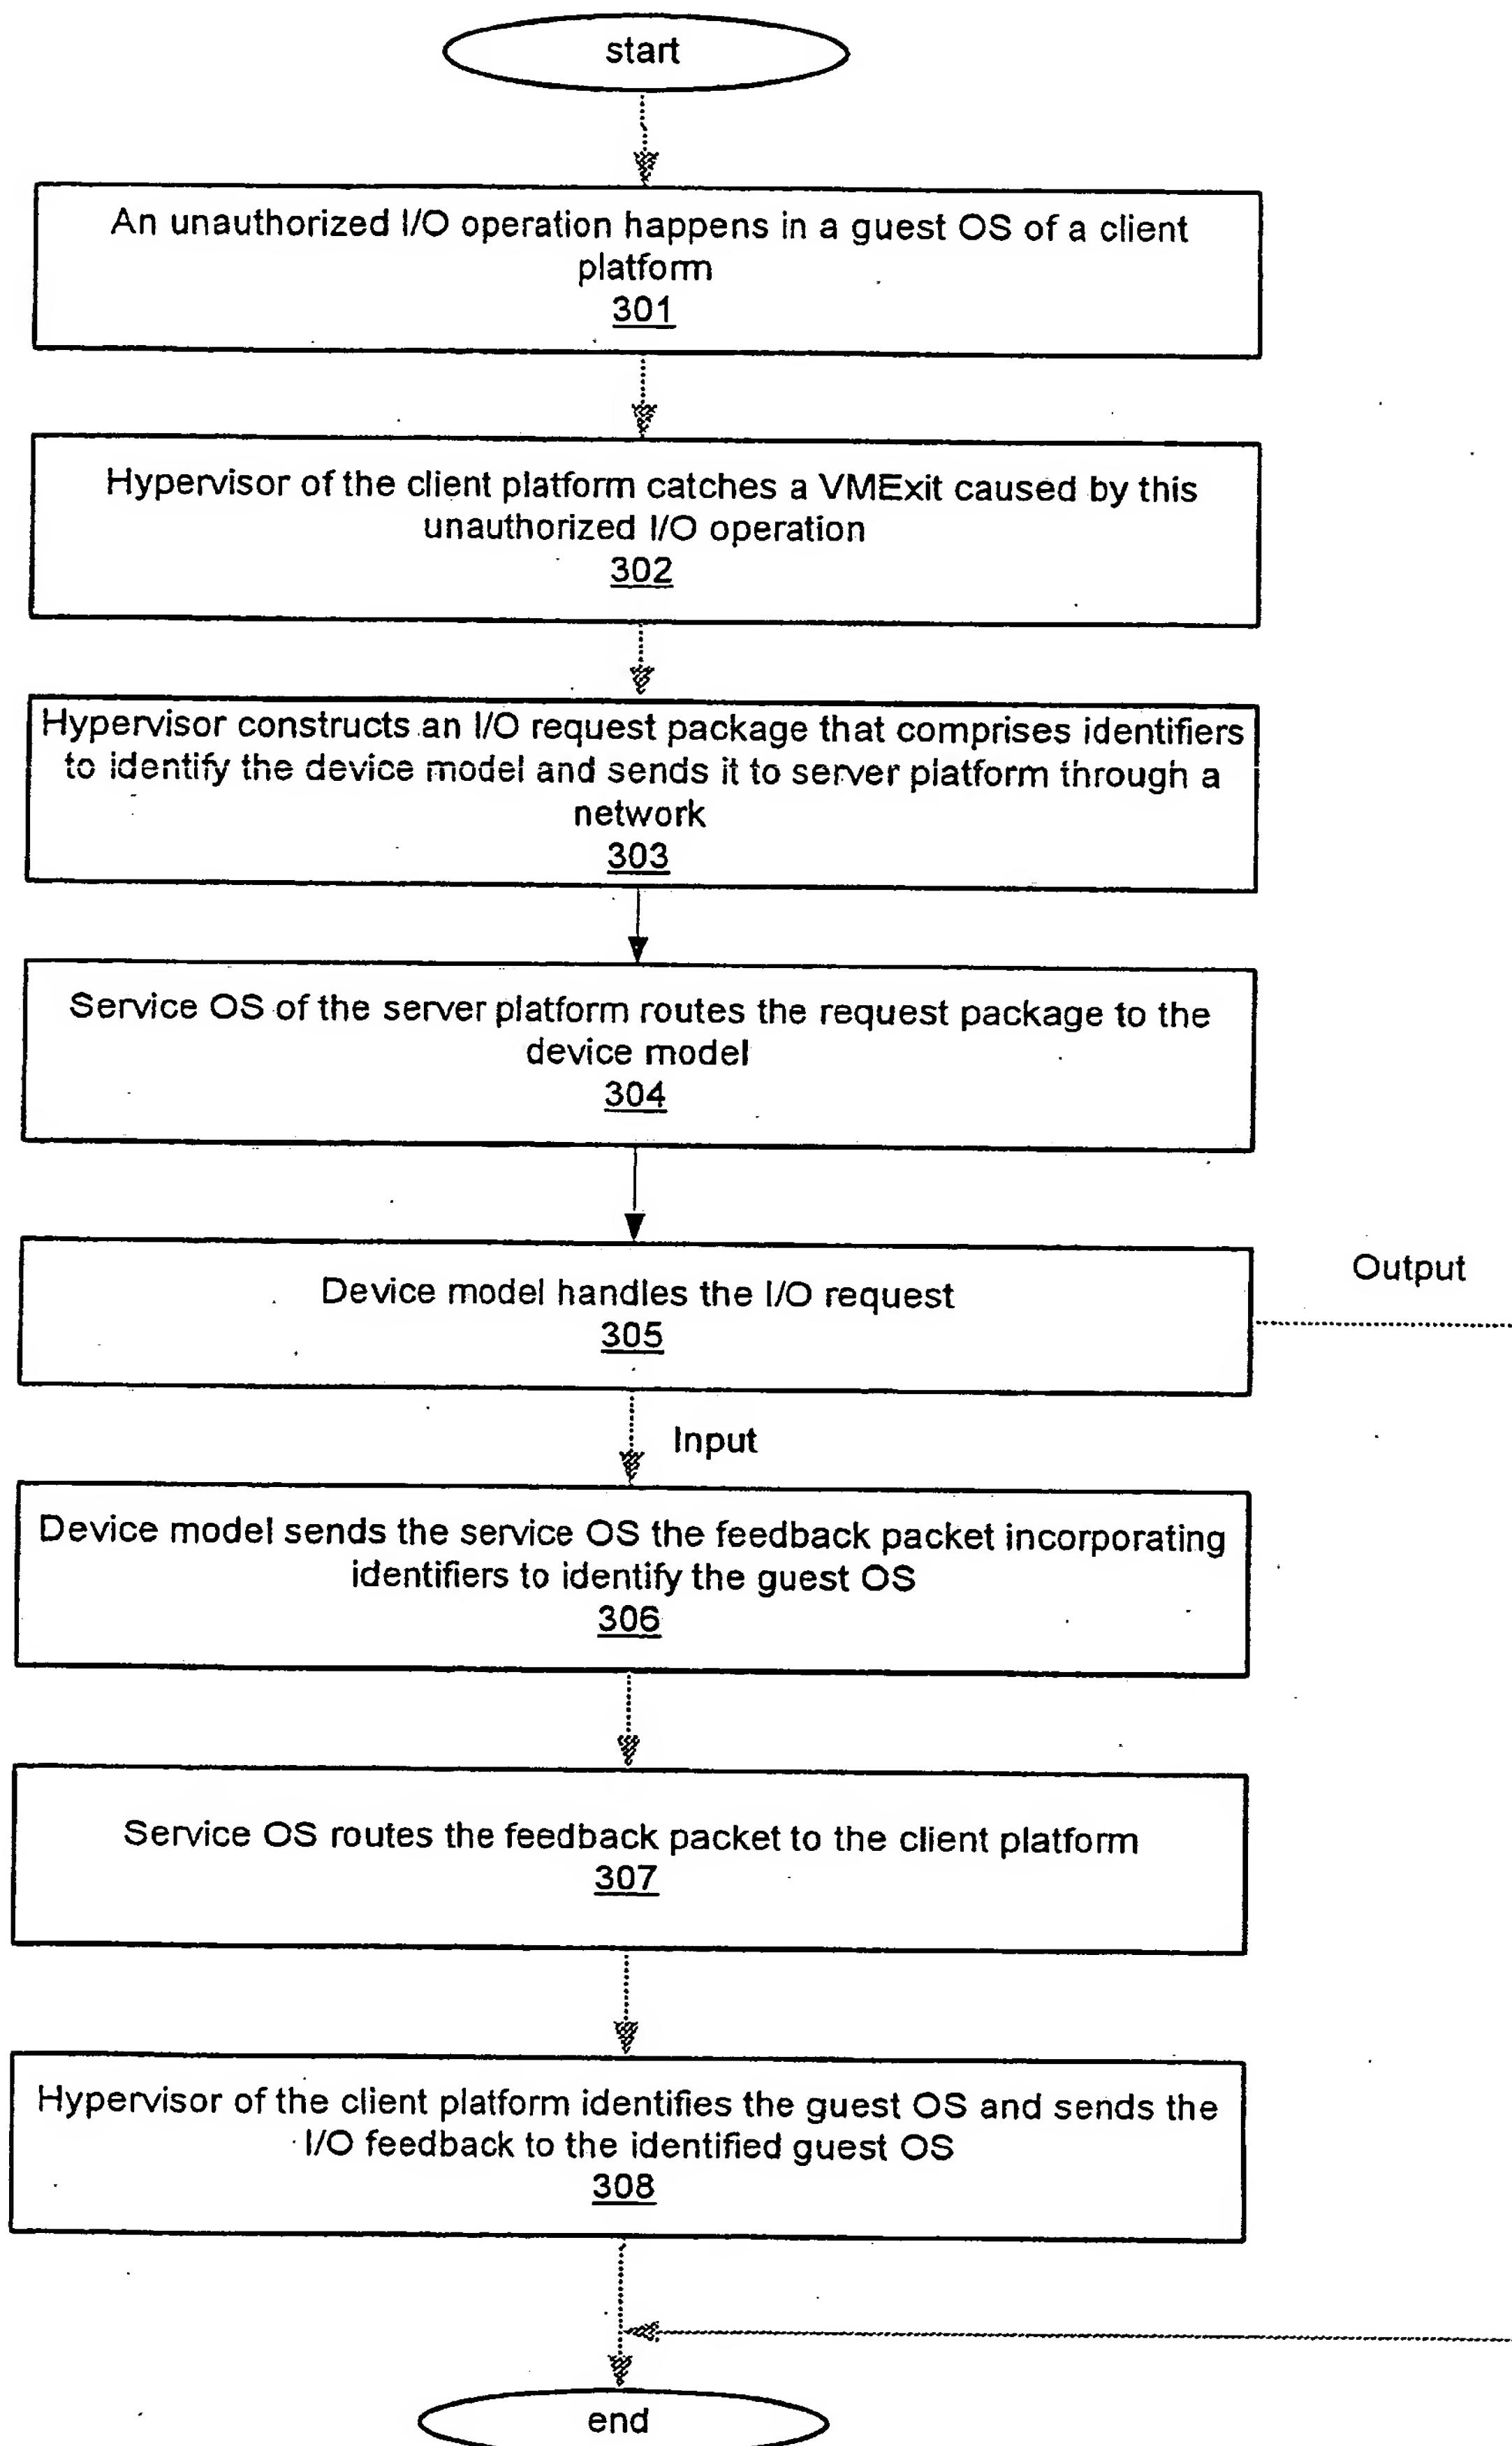
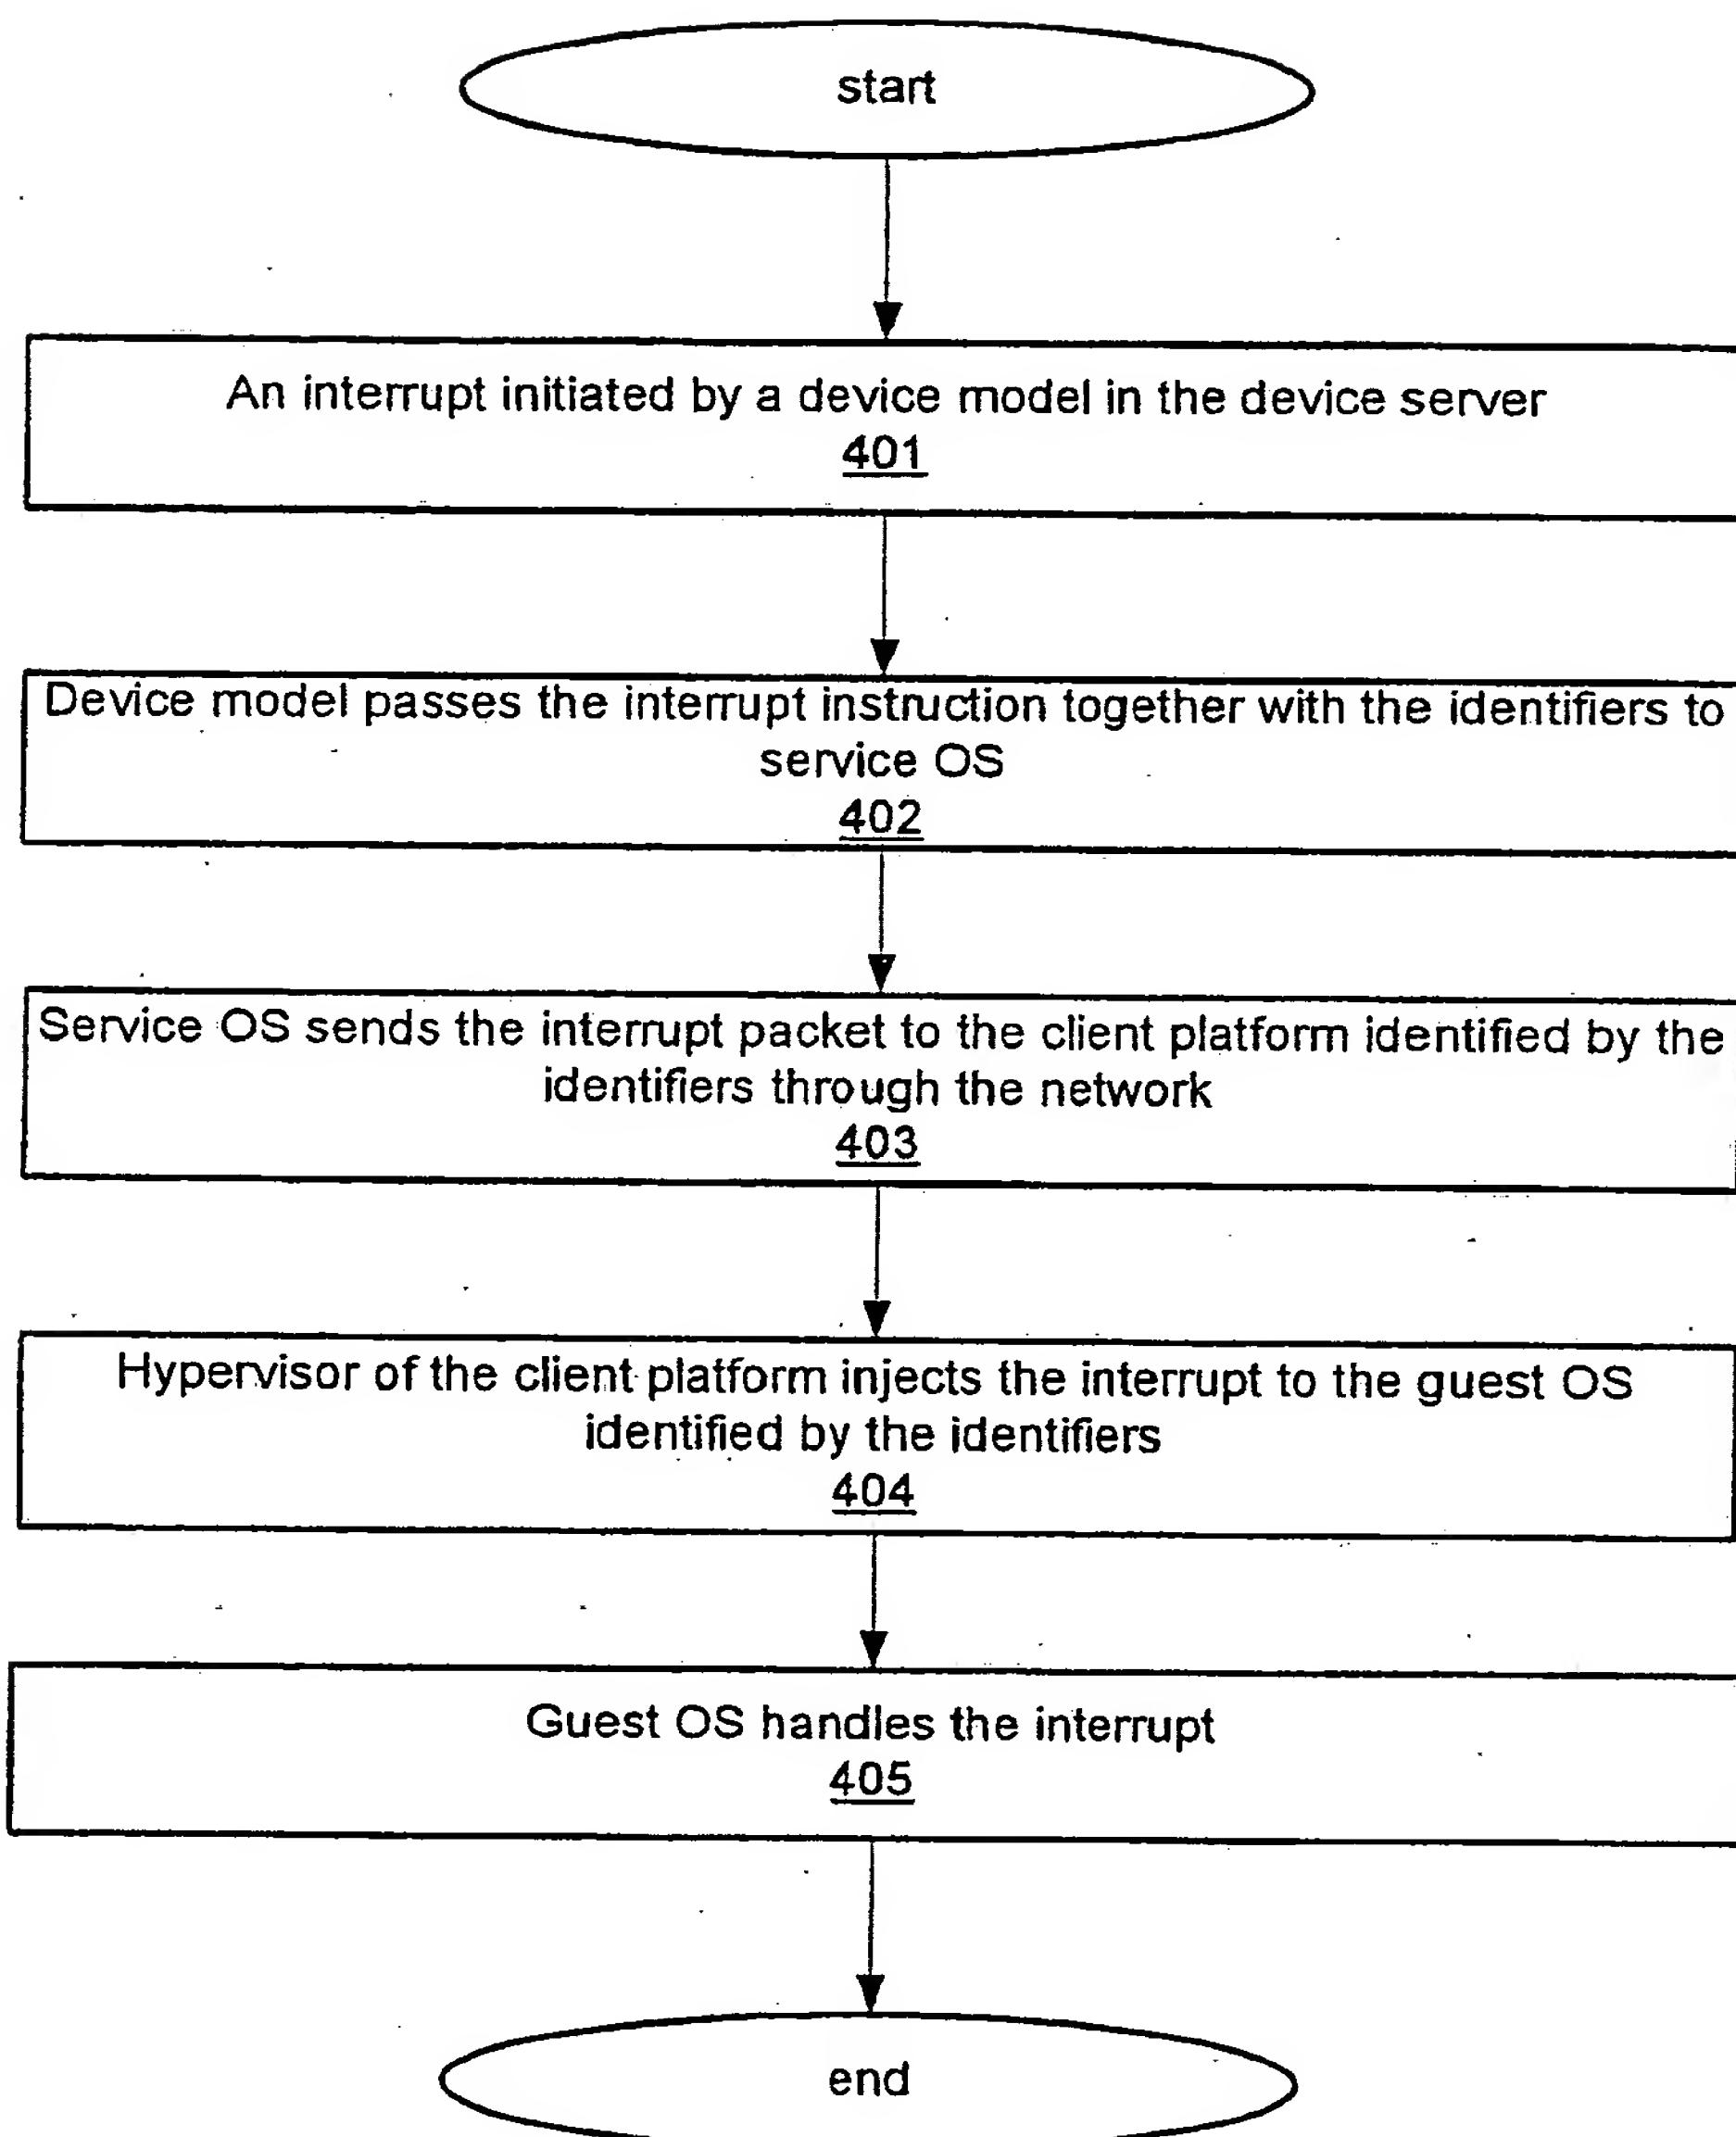
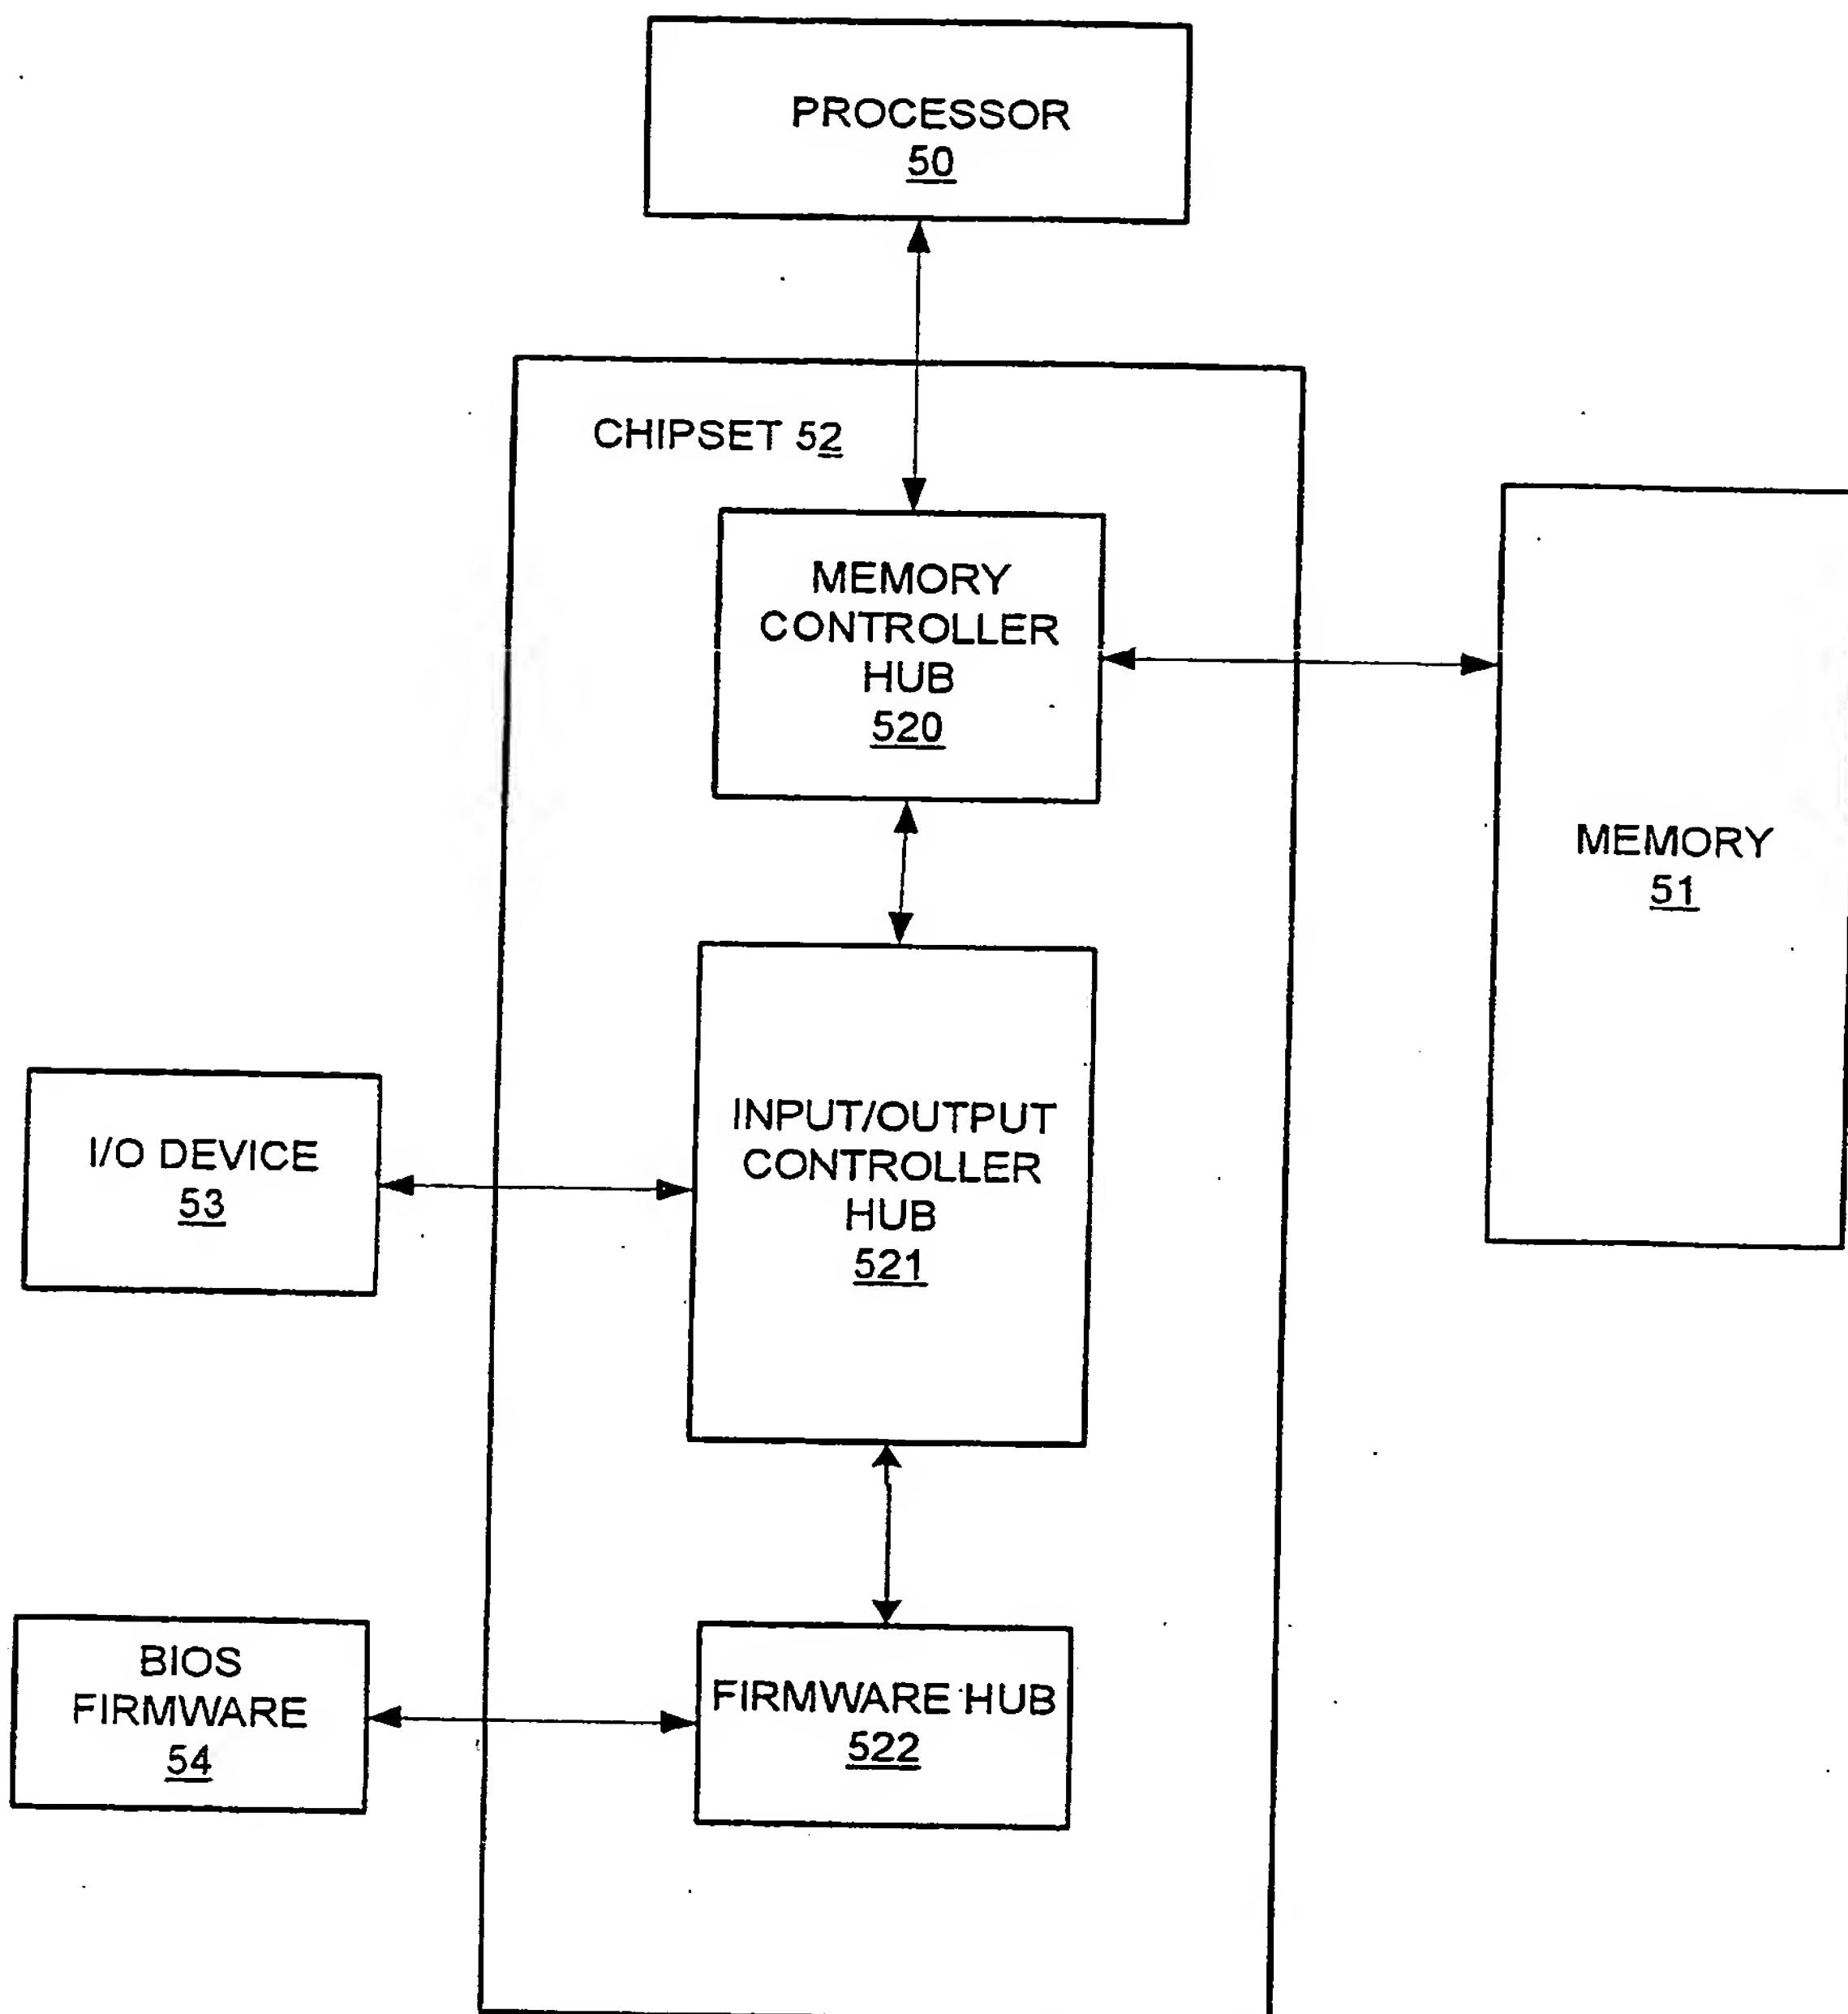
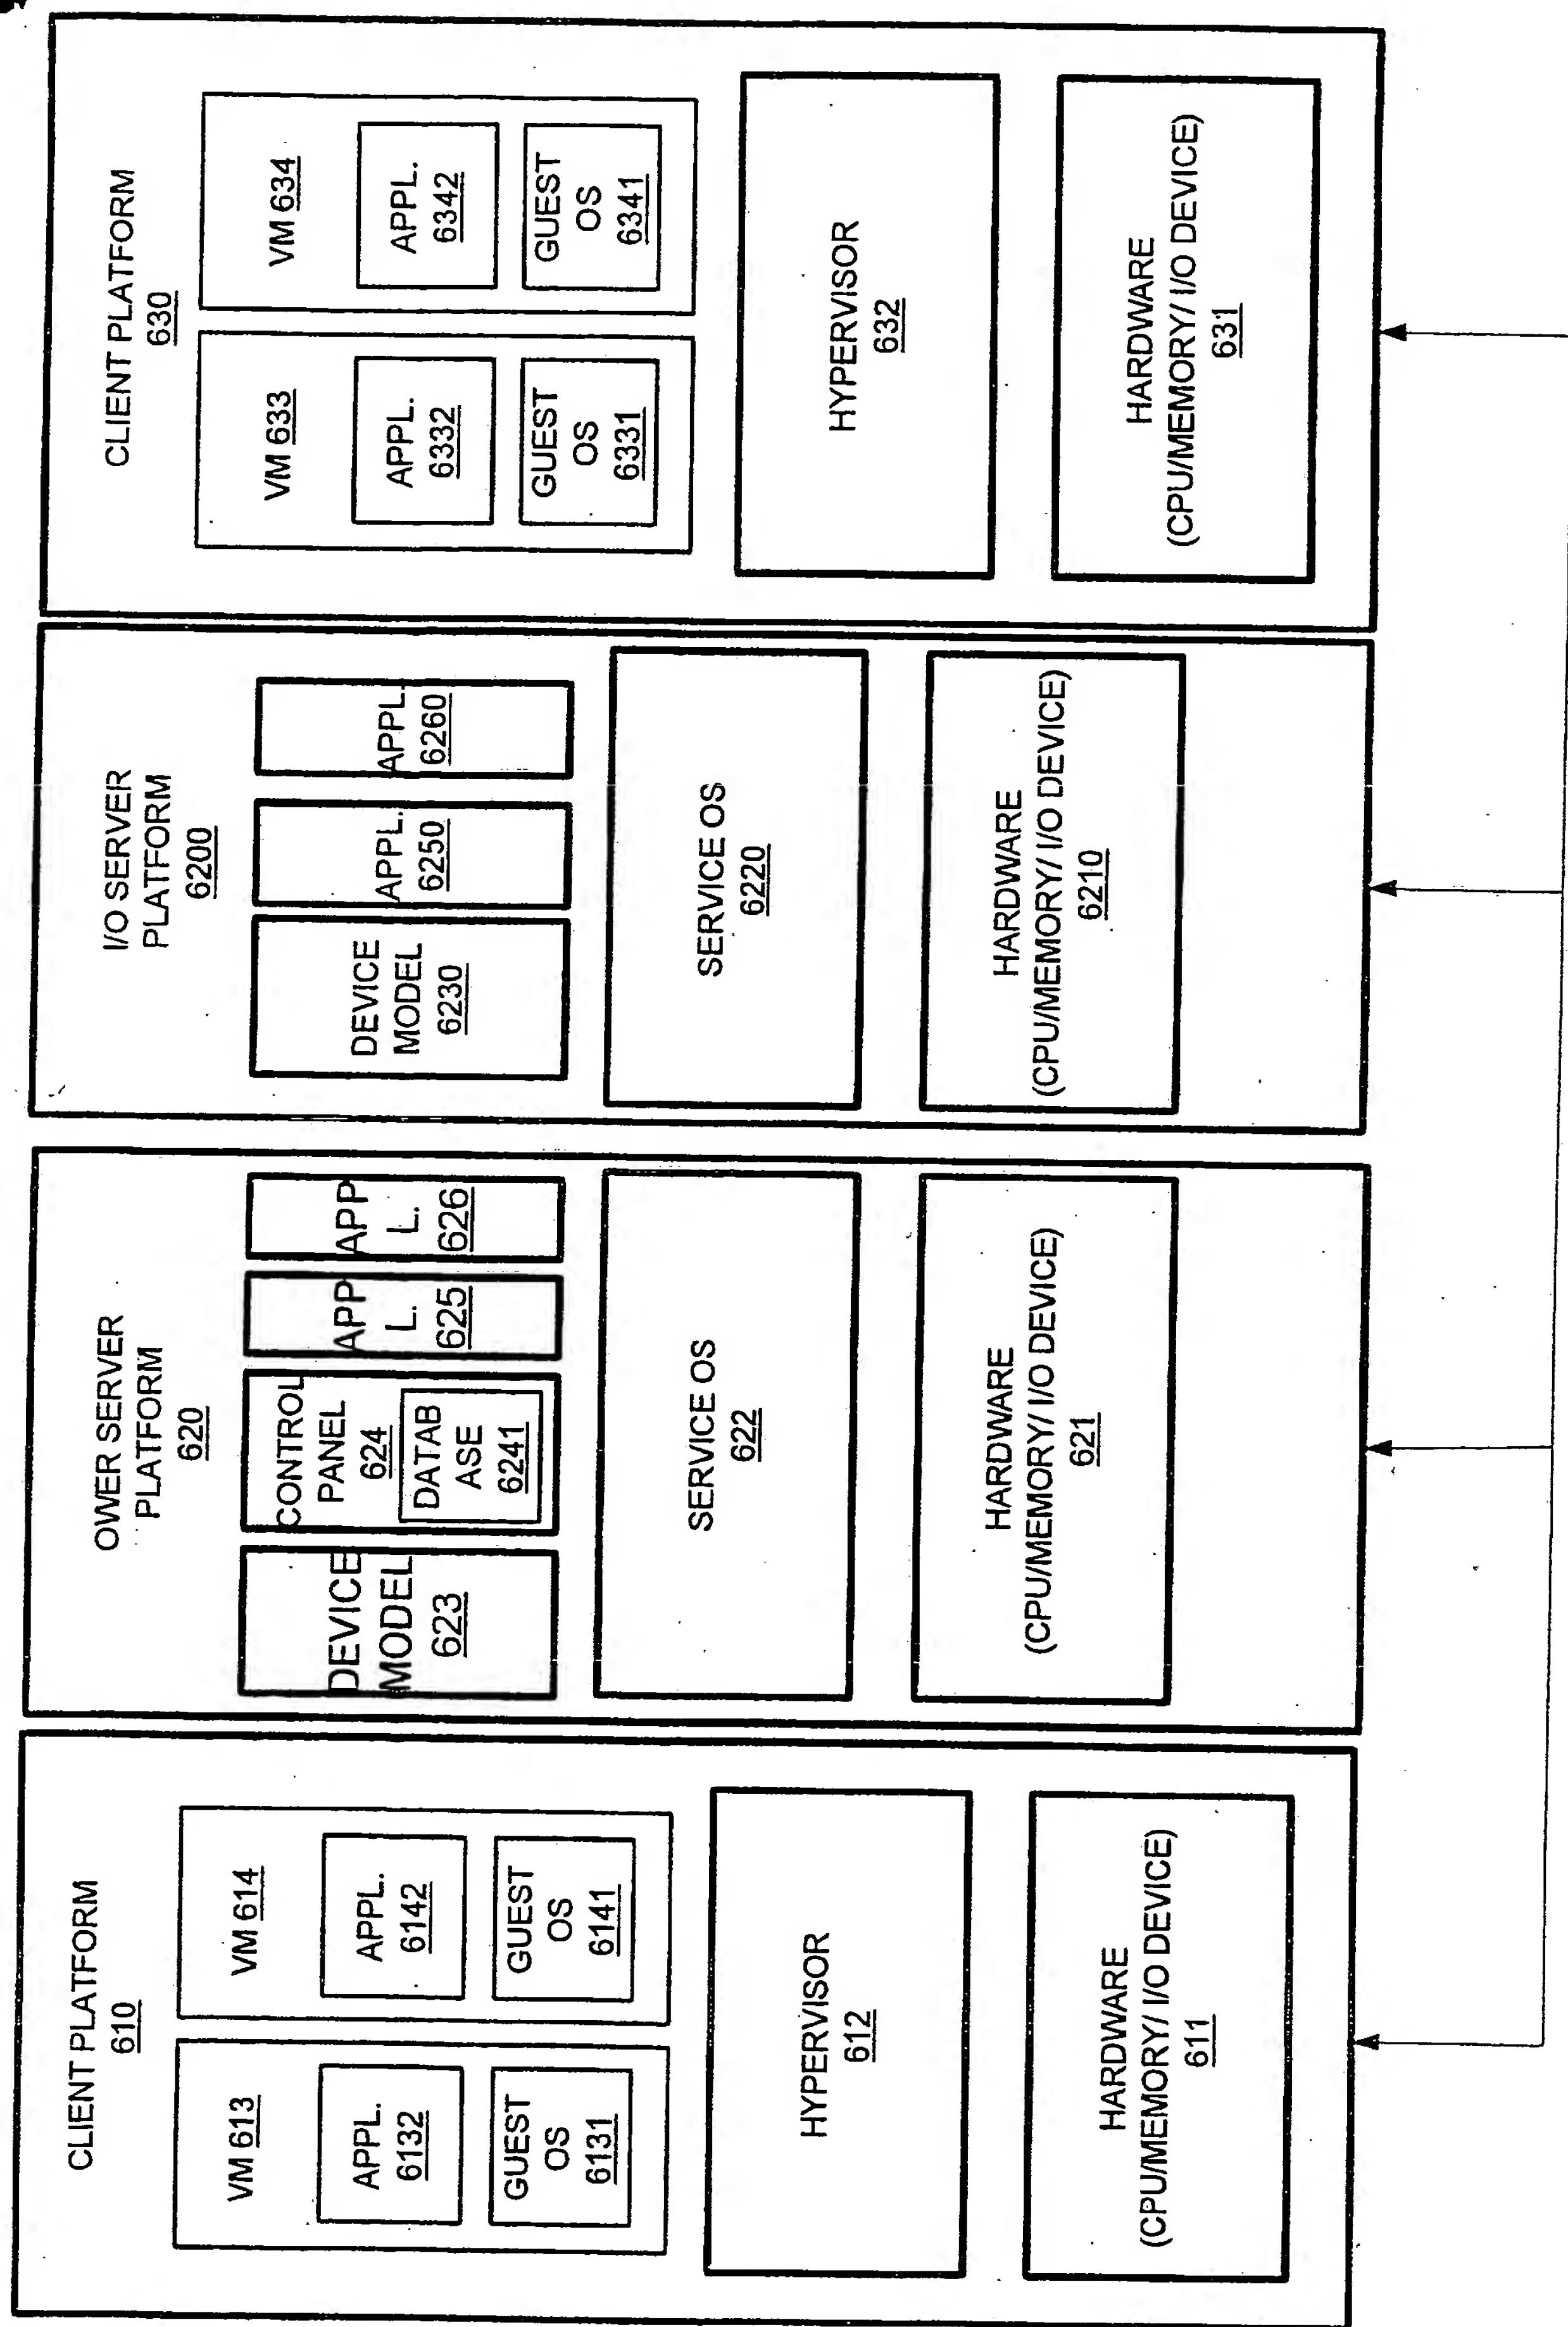


FIG. 3

4/6

**FIG. 4**

**FIG. 5**



NETWORK 40

FIG. 6